Troubleshooting the Cisco Intercloud Fabric Components

Issues with the Intercloud Fabric Virtual Supervisor Module

This section includes symptoms, possible causes, and solutions for issues associated with the Cisco Intercloud Fabric virtual supervisor module (VSM).

Symptom: In the output of the show module command, the Intercloud Fabric Switch modules are not shown online.

Possible Cause: The Intercloud Fabric Switch is not registered with the Cisco Intercloud Fabric VSM due to one of the following reasons:

- The Intercloud Fabric Switch management VLAN is missing from the Intercloud Fabric cloud trunk port profiles. The port profiles are named IcfCloud_Name_ICS_Trunk_Tunnel, IcfCloud_Name_ICX_Trunk_Tunnel, and IcfCloud_Name_ICX_Trunk_Tunnel_Enterprise.
- The Intercloud Fabric Switch management VLAN is not identified as the system VLAN for the Intercloud Fabric cloud trunk port profiles. The port profiles are named IcfCloud_Name_ICS_Trunk_Tunnel, IcfCloud_Name_ICX_Trunk_Tunnel, and IcfCloud_Name_ICX_Trunk_Tunnel_Enterprise.

Verification and Solution: Manually add the VLAN configuration interface in ICFC.

Symptom: There is no traffic from the private VM to the cloud VM.

Possible Cause: The cloud VM VLAN is missing from the Intercloud Fabric cloud trunk port profiles. The port profiles are named IcfCloud_Name_ICS_Trunk_Tunnel, IcfCloud_Name_ICX_Trunk_Tunnel, and IcfCloud_Name_ICX_Trunk_Tunnel_Enterprise.

Verification and Solution: Manually add the VLANs via the VSM configuration interface in ICFC.

Issues with Intercloud Fabric Controller

This section includes symptoms, possible causes, and solutions for issues associated with the Intercloud Fabric Controller (ICFC).

Symptom: The UI task status displays a failure message.

Possible Cause: The possible causes vary.

Verification and Solution: Review the resource manager logs in the `/var/log/resource-mgr` folder. The `svc_res_dme.log` log file is for UI-to-Resource Manager communication. The `svc_res_cloudproviderAG.log` log file is for AG-to-CPM communication.

Symptom: The access tunnel fails.

Possible Cause: The possible causes vary.

Verification and Solution: Review the logs on the Intercloud Fabric Switch.

Symptom: The Intercloud image fails to import into ICFC.

Possible Cause: ICFC encountered an error.

Verification and Solution: In ICFC, review the **Recent Jobs** list for the status and an error message if the import operation does not complete within a reasonable amount of time. The **Task** line item lists the last error encountered, if any, and the number of retries attempted.

Symptom: After restarting ICFC services, the job for importing an Intercloud image is displayed as running, but the job does not progress.

Possible Cause: If ICFC services are restarted while an import operation is running, the task does not resume.

Verification and Solution: Cancel the job in Intercloud Fabric and import the image again.

Symptom: Intercloud Fabric cannot validate Amazon Web Services (AWS) credentials.

Possible Cause:

- Not all of the required ports are open.
- The DNS server is not reachable from ICFC.
- The NTP server is not reachable from ICFC.

Verification and Solution:

1. Make sure that all required ports are open:
    - TCP ports—22, 443, 3389, 6644, and 6646
    - UDP ports—6644 and 6646
2. Make sure that the DNS service is configured and the DNS server is reachable from ICFC.
3. Make sure that the NTP service is configured and the NTP server is reachable from ICFC.

Symptom: Integrated Gateway is unreachable from ICFC.

Possible Cause: After restarting ICFC or the Intercloud Fabric VM, the Integrated Gateway is unreachable.

Verification and Solution:

1. SSH to ICFC and enter the following details:

```
ICFC# connect local-mgmt
ICFC(local-mgmt)# service restart
```

2. In ICFC, check if the Integrated Gateway state changes to reachable/running.

Note If the Integrated Gateway state does not change to reachable/running, change an Integrated Gateway configuration parameter, such as the description. This changes the Integrated Gateway state to reachable/running by attempting a new connection.

Symptom: An Intercloud Fabric link is deployed but the configuration status remains "Applying" and the tunnel status shows "not operational". Similarly, when ICS is successfully deployed and obtains a private IP address (with a configuration status "Applying"), the configuration status indicates that ICS is not reachable from ICFC.

Possible Cause: ICFC must be behind NAT for ICFC-to-ICS communication.

Verification and Solution: Modify or delete the route on ICS. ICFC should be able to reach ICS via provider private gateway.

Symptom: ICFC fails to respond or come up.

Possible Cause: Multiple situations can cause this problem.

Verification and Solution:

1. Using SSH, connect to the Intercloud Fabric VM CLI as an administrator.
2. Enter option **11** to display the container details for ICFC information and current status.
3. Enter option **15** to launch the ICFC console for further detailed ICFC information.

Note For more information about Intercloud Fabric VM CLI console, see the *Cisco Intercloud Fabric Administrator Guide*, section "Using Cisco Intercloud Fabric VM CLI Commands."

Symptom: The ICFC container status shows that it is running but fails to respond to queries.

Possible Cause: Multiple situations can cause this problem.

Verification and Solution:

1. Using SSH, connect to the Intercloud Fabric VM CLI as an administrator.
2. Enter option **15** to launch the ICFC console.
3. Run the 'connect local-mgmt' command, followed by the 'service status' command, on the ICFC CLI to determine the ICFC service status.

Note For more information about Intercloud Fabric VM CLI console, see the *Cisco Intercloud Fabric Administrator Guide*, section "Using Cisco Intercloud Fabric VM CLI Commands."

Issues with the Intercloud Fabric Switch

This section includes symptoms, possible causes, and solutions for issues associated with the Intercloud Fabric Switch (ICS). The symptoms include the following:

Symptom: A site-to-site tunnel does not come up.

Possible Cause: For TCP and UDP tunnel type:

- Port 6644 on ICX is not open. Both UDP and TCP tunnel types need port 6644 open for the tunnel to be established.

Verification and Solution: To open port 6644 behind the firewall, complete the following steps:

1. Verify that the tunnel received the configuration parameters (which describe the site-to-site tunnel type) and the tunnel state by entering the following command:

   ```
   show intercloud tunnel config
   ```

2. Verify that the tunnel state is up by entering the following command on ICX and ICS:

   ```
   show intercloud tunnel brief
   ```

   For HTTPS tunnel on ICX, continue with the following steps.

3. Display the counter related to the tunnel establishment failure by entering the following commands:

   ```
   show intercloud tunnel statistics

   show systm internal tunnel-manager dp
   ```

4. To collect the log information, enter the following command:

   ```
   show tech-support
   ```

   Note An intermittent network issue may cause the site-to-site tunnel from coming up. A reboot of the site-to-site link from the Intercloud Fabric UI may fix this issue. In result, a high availability failover may be triggered if high availability is configured. If there is no high availability, it will result in all cloud VMs to lose connectivity with the cloud provider for the duration of the site-to-site reboot operation.

5. Repeat step 3 on ICS and collect the statistics and log and forward to Cisco Customer Support.

Symptom: Access tunnel does not come up.

Possible Cause:

Verification and Solution:

1. Verify the tunnel has received the configuration parameters and the tunnel state status by entering the following command:

   ```
   show intercloud tunnel config
   ```

2. Verify that the tunnel state is up by entering the following command on ICX and ICS:

   ```
   show intercloud tunnel brief
   ```

   If the tunnel has not received tunnel configuration parameters, the first command will return blank entries for tunnel parameter for access configuration.

3. Reboot the site-to-site link from the Intercloud Fabric UI to force ICS to receive access tunnel management parameters toward the cloud VM.

4. If ICS has received the access tunnel configuration parameters, review if any ICS cores are present by entering the following command:

   ```
   show cores > bootflash/cores.log
   ```

5. Collect the ICS logs by entering the following command:

   ```
   show tech-support > bootflash/logs
   ```

6. Copy any cores and logs and forward to Cisco Customer Support.

7. Reboot the cloud VM.

   Note A reboot of the cloud VM from the Intercloud Fabric UI will result in a cleaned cloud VM. ICS will then resume connecting with the cloud VM and open an access tunnel after the reboot.

8. If the issue remains, collect any remaining cores and logs and contact Cisco Customer Support for further troubleshooting.

Symptom: A cloud VM cannot reach another cloud VM on the same VLAN.

Possible Cause:

Verification and Solution:

1. Verify the tunnel has received the configuration parameters by entering the following command:
   ```
   show intercloud tunnel config
   ```
   If the tunnel has not received tunnel configuration parameters, this command will return blank entries for tunnel parameters for access configuration.
2. Reboot the site-to-site link from the Intercloud Fabric UI to force ICS to receive access tunnel management parameters toward the cloud VM.
3. If ICS has received the access tunnel configuration parameters, review if any cores are present by entering the following command:
   ```
   show cores > bootflash/cores.log
   ```
4. Collect the ICS logs by entering the following command:
   ```
   show tech-support > bootflash/logs
   ```
5. Copy any cores and logs and forward to Cisco Customer Support.
6. Reboot both cloud VMs.

   Note A reboot of the cloud VM from the Intercloud Fabric UI will result in a cleaned cloud VM. ICS will then resume connecting with the cloud VM and open an access tunnel after the reboot.
7. Review the tunnel access state for both cloud VMs. If the tunnels are up, connectivity should be established between the two cloud VMs by entering the following command:
   ```
   show intercloud tunnel brief
   ```
8. If the issue remains, contact Cisco Customer Support after collecting any cores and logs by entering the following command:
   ```
   show tech-support > bootflash/logs
   ```

Symptom: Packets from the cloud provider cannot reach the cloud VM on the same VLAN.

Possible Cause: One of the following may be dropping packets:

- ICX
- ICS
- End VM

Verification and Solution: To debug the packets lost, complete the following steps for either ICX, ICS (for HTTPS tunnel), or ICS (in case of UDP tunnel):

- ICX

1. Verify the tunnel has received the configuration parameters and the tunnel state status by entering the following command:
   ```
   show intercloud tunnel config
   ```
   This command verifies the type of site-to-site tunnel operation.

2. Verify the tunnel is up by entering the following command:
   ```
   show intercloud tunnel brief
   ```
   If the tunnel is up, proceed to step 7.

3. Verify the VEM state for the presence of cloud port and VLAN membership for the port by entering the following commands:
   ```
   vemcmd show port

   vemcmd show port vlans
   ```

4. Review the VEM packet counters on ICX by entering the following commands:
   ```
   vemcmd show stats

   vemcmd show port-drops ingress/egress
   ```
   If the drop count for the port associated with the cloud facing port is going up, packets may be getting dropped. This may be due to the port, not being configured for VLAN membership, getting extended into the cloud.

5. Determine if the cloud facing port is configured properly for VLAN membership by entering the following command:
   ```
   vemcmd show port vlans
   ```

6. For HTTP site-to-site tunnel, review the ICX tunnel counters and verify if the drop counter for the tunneled packets is going up by entering the following commands:
   ```
   show intercloud tunnel statistics

   show system internal tunnel-manager dp
   ```
   If the drop counter for the tunneled packets is going up, this may be due to an installation error. Contact Cisco Customer Support.

7. If packets are successfully moving out of ICX, check the following:
   - For ICS (for HTTPS tunnel):
     ```
     show system internal tunnel-manager dp
     ```
   - For ICS (in case of UDP tunnel):
     ```
     show intercloud tunnel statistics
     ```

Note Review the packet drop counters. If the cloud VM port counters for data transmitted and received are moving up, the packets are reaching the cloud VM. In this case, log in to the cloud VM and review the packet flow on the cloud VM.

Symptom: Packets from the cloud VM cannot reach the cloud provider.

Note This issue assumes that the cloud VM access tunnel is up and working. If the cloud VM and the cloud provider VM is not the same VLAN, you can use one to verify the packet path from cloud VM to ICS. The Routing Service is responsible for packet switching from one VLAN to another VLAN. Refer to Issues with the Routing Service section for Routing Service troubleshooting operations.

Possible Cause: The possible causes include the following:

- The cloud VM is not configured with routes pointing toward the cloud provider.
- ICS is dropping packets.
- ICX is dropping packets.
- End VM is dropping packets.

Verification and Solution: Debug the packets lost by performing the following steps for ICS and ICX:

- For ICS, verify the route table entries on the cloud VM point toward the cloud provider VM:

1 Review the route table from the cloud VM by entering the following command:

```
show intercloud vm-name <Name of the VM> system-log
```

If the route table entries are accurate, review the packet counters on ICS to verify the packets are reaching ICS.

2 Review the interface counters by entering the following command:

```
show intercloud tunnel statistics
```

If the end VM is not on the same VLAN, but the packet counter for the access tunnel are function properly, use the Routing Service debug operations to verify that the packets get the VLAN switched. If the end VM is on the same VLAN as the cloud VLAN, continue to step 3.

3 Review the site-to-site tunnel type by entering the following command:

```
show intercloud tunnel config
```

4 Verify the site-to-site tunnel is up by entering the following command:

```
show intercloud tunnel brief
```

5 Review the HTTP site-to-site tunnel counters by entering the following commands:

```
show intercloud tunnel statistics
```

```
show system internal tunnel-manager dp
```

Note If the drop counter for the tunneled packets is going up, this may be due to an installation error. Contact Cisco Customer Support.

If packets are successfully moving out of ICS, proceed to the ICS procedure.

6 Review ICS (for HTTPS tunnel) by entering the following command:

```
show intercloud tunnel statistics
```

Review the packet drop counters. If the site-to-stie port counters for data transmitted and received are moving up, the packets are reaching ICX. In this case, log in to ICX and review the packet flow.

7 If packets are successfully moving out of ICX, check the following:

- For ICS (for HTTPS tunnel):
  ```
  show system internal tunnel-manager dp
  ```
- For ICS (in case of UDP tunnel):
  ```
  show intercloud tunnel statistics
  ```

Note Review the packet drop counters. If the cloud VM port counters for data transmitted and received are moving up, the packets are reaching the cloud VM. In this case, log in to the cloud VM and review the packet flow on the cloud VM.

- Perform the following steps for ICX:

1 Verify the site-to-site tunnel is up by entering the following command:

```
show intercloud tunnel brief
```

2 Review the HTTP site-to-site tunnel counters on ICX by entering the following commands:

```
show intercloud tunnel statistics

show system internal tunnel-manager dp
```

Note If the drop counter for the tunneled packets is going up, this may be due to an installation error. Contact Cisco Customer Support.

If packets are successfully moving out of ICX, review the end VM.

3 Review ICX (in case of UDP tunnel) by entering the following command:

```
show intercloud tunnel statistics
```

Note Review the packet drop counters. If the cloud VM port counters for data transmitted and received are moving up, the packets are reaching the cloud provider VM. In this case, log in to the cloud provider VM and review the packet flow.

Symptom: ICS fails to power up and AWS displays a "401 Unauthorized" error.

Possible Cause: The credentials were changed or the access key was deleted on AWS.

Verification and Solution: Enter the correct credentials for the provider account. For information about entering the credentials, see the Amazon Web Services documentation.

Symptom: ICS does not appear under "show module" for the registered cVSM.

Possible Cause:

- VSM reachability from ICS over the management interface is not working.
- One of the ports on ICX is not set to trunk mode.
- The host port is not set to promiscuous mode. As a result, the ICX host does not receive all packets from ICS.

Verification and Solution:

1. When using different VLANs for ICS, either add a static route for that subnet or make sure the subnet can be reached through the default gateway. To configure this, choose **Intercloud > Network > IcfVSM**, highlight the VSM, and click **Configure Route**.
2. Log in to the VMware vCenter client. Right-click the ICX VM and choose **Edit Settings**. Three network adapters appear. Verify that the correct NIC (the NIC facing the cloud side) is set to "trunk mode".
3. Select the ICX host and click the **Configuration** tab. Select **Networking**, choose the vSwitch responsible for the port (physical adapter) serving the ICX VM, and click **Properties**. In the **Ports** tab, click **Edit** and navigate to **Security**. Set the following to "accept mode:"
    - Promiscuous mode
    - MAC address change
    - Forged transmits

Symptom: Packets arriving at the ICX from the private cloud are not sent to the site-to-site tunnel to the ICS.

Possible Cause: The possible causes vary.

Verification and Solution: To debug packets lost, enter the following commands:

```
show intercloud tunnel brief
show intercloud tunnel config

vemcmd show port
vemcmd show port vlans
vemcmd show stats
vemcmd show port-drops ingress/egress
show system internal tunnel-manager dp
show system internal event-log tunnel manager
```

Symptom: Unknown unicast packets, with MAC addresses that are not in the MAC address pool, are not forwarded over the site-to-site tunnel.

Possible Cause: On the ICX, unknown unicast packets forwarded from the private cloud to the cloud over the site-to-site tunnel are dropped if their destination MAC address is not a part of the MAC address pool for the Intercloud Fabric link.

Verification and Solution: Use the following command to display the feature status for the blocking of unknown unicast MAC addresses that are not in the MAC address pool:

```
vemcmd show macpool veth1-0
```

This command applies only to the ICX. The feature is on by default. Contact customer support if you need to disable this feature.

Symptom: The overlay address on the cloud VM is not configured.

Possible Cause: The subagent process did not receive or push the IP address configuration from the ICS to the cloud VM, but the control channel between the ICS and the cloud VM is verified as up and working.

Verification and Solution: Open a session to the ICS and enter the following command:

```
show system internal event-log tunnel-mgr
```

Issues with the Intercloud Fabric Routing Service

This section includes symptoms, possible causes, and solutions for issues associated with the Routing Service.

Issues Enabling and Configuring Intercloud Fabric Routing Service

Symptom: Creating the Routing Service fails.

Possible Cause:

- The transport network does not exist.
- There are not enough IP addresses in the transport network.
- There are not enough IP addresses for the management IP addresses.
- The ICFC tunnel is down.
- The Intercloud Fabric Switch cannot be reached from ICFC.
- There are not enough MAC addresses in ICFC.
- ICFC cannot be reached.

Verification and Solution:

1. Create a transport network, redeploy the link, and review the service request logs for warnings:

   ```
   Routing Service creation aborted. No transport network found.
   ```

2. Allocate more IP addresses to the transport network, redeploy the link, and review the service request logs for warnings:

   ```
   Routing Service creation aborted. No ips in transport netowrk <network name>.
   ```

3. Allocate more IP addresses to the management network and review the service request logs for warnings:

   ```
   Routing Service management IP allocation failed pool ID <oid>.
   ```

Use this procedure if ICFC tunnel is down:

1. Log in to cVSM and make sure the tunnel endpoints are registered.
2. Log in to the Intercloud Fabric Switch and confirm that the tunnel is operational.
3. From the Intercloud Fabric Switch, confirm that the Routing Service service is running.
4. From the Intercloud Fabric Switch, confirm that the data is correct for the port configuration of the Routing Service trunk.
5. Confirm that the port profile configurations for the cVSM are correct.

   Review the service request logs for warnings:

   ```
   Failed to create the Routing Service instance <Routing Service IP address> for icfCloud.
   This is primarily due to a network connectivity issue between ICFC and ICS.
   ```

Use this procedure if the Intercloud Fabric Switch cannot be reached:

1. Ping the Intercloud Fabric Switch IP address and confirm that the configurations are correct.
2. Log in to cVSM and confirm that the tunnel endpoints are registered.

3 Confirm that the port profile configuration for cVSM are correct.

Review the service request logs for warnings:

```
Failed to create the Routing Service instance <Routing Service IP address> for icfCloud.
This is primarily due to a network connectivity issue between ICFC and ICS.
```

If there are not enough MAC pool addresses, review the service request logs for warnings:

```
If ICFC cannot be reached, make sure the information provided during the OVA installation
 is accurate.
If ICFC cannot be reached due to invalid information during installation, reinstall the
 OVA with the correct information.
```

Symptom: Configuring the Routing Service fails.

Possible Cause: The Routing Service login fails.

Verification and Solution: Review the service request logs for warnings.

Symptom: Creating a Routing Service fails because no transport Network is present in Intercloud Fabric.

```
Routing Service creation aborted. No transport network found.
```

Possible Cause: The presence of a transport Network is required to create a Routing Service.

Verification and Solution:

1 Review the Intercloud Fabric service request logs.
2 Create the transport Network in Intercloud Fabric.
3 Disable the Routing Service on the Intercloud Fabric cloud.
4 Once successfully disabled, reenable the Routing Service on the Intercloud Fabric cloud.

Symptom: Creating a Routing Service fails because the transport Network does not have a free IP address.

```
Routing Service creation aborted. No ips in transport network <network name>.
```

Possible Cause: Routing Service creation fails when the transport Network is different from the management Network.

Verification and Solution:

1 Review the Intercloud Fabric service request logs.
2 Allocate at least one more IP address to the transport Network.
3 Disable the Routing Service on the Intercloud Fabric cloud.
4 Once successfully disabled, reenable the Routing Service on the Intercloud Fabric cloud.

Symptom: Creating a Routing Service fails because the management Network does not have one (or two, if in high availability mode) free IP address(es).

```
Routing Service management IP address allocation from <network name> failed.
```

Possible Cause: Routing Service creation fails when the availability of one (or two) free IP address(es) on the management Network is not present.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Allocate at least one (or two, if in high availability mode) more IP address(es) to the management Network.
3. Disable the Routing Service on the Intercloud Fabric cloud.
4. Once successfully disabled, reenable the Routing Service on the Intercloud Fabric cloud.

Issues Disabling Intercloud Fabric Routing Service

Symptom: Removing the Routing Service IP address fails if the Routing Service is removed while a Cisco Intercloud Fabric link is being deleted.

Possible Cause:

- The management IP address removal fails.
- The interface IP address removal fails.

Verification and Solution:

1. Delete the Routing Service-related object so that the Cisco Intercloud Fabric link can be deleted. Review the service request logs for warnings:

   ```
   Failed to release Mgmt IP <ips> for IG Instance <Routing Service name> (error
   message).
   ```

2. If deleting the Cisco Intercloud Fabric link does not succeed, review the service request logs for warnings:

   ```
   Failed to release Ips <ipAddresses> allocated to Routing Service networkInterface
   <icRouterInterface. name>.
   ```

Symptom: Deleting the Routing Service fails.

Possible Cause:

- ICFC displays an error.
- ICFC times out.

Verification and Solution:

1. Confirm that the ICFC IP address can be reached and review the service request logs for warnings.
2. Confirm that the Intercloud Fabric Switch container exists and review the service request logs for warnings.

Symptom: The Routing Service interface IP address release fails.

Possible Cause: The allocation failure occurs because of one, or more, of the following interface types:

- Management Network interface
- Transport Network interface
- Data only Network interface

Verification and Solution:

1. Review the Intercloud Fabric service request logs.

   ```
    Failed to release Mgmt IP <ips> for IG Instance <Routing Service name> (error
   message).
   ```

2. Review the workflow error logs for further debugging.

   ```
   Failed to release Mgmt IP <ips> for IG Instance <Routing Service name> (error
   message).
   ```

Issues Connecting Intercloud Fabric Routing Service to Networks

Symptom: Creating an interface fails when connecting the Routing Service to networks.

Possible Cause:

- There are insufficient IP addresses for one or more networks.
- Creating an ICFC network interface fails.
- ICFC times out.
- The Routing Service does not appear in the system or is in a failed state.
- The Intercloud Fabric Switch or the integrated gateway cannot be reached.

Verification and Solution:

1. Review the service request logs for warnings.
2. If ICFC times out, confirm that the Routing Service interface is properly configured and that the ICFC IP address can be reached. In addtion, review the service request logs for warnings.

Symptom: No interfaces are created on the Routing Service when connecting to networks.

Possible Cause:

- No data Layer 3 network (other than management and transport) is associated with the VDC.
- The interface already exists.
- The Routing Service is not in the system or is in a failed state.

Verification and Solution: Review the service request logs for warnings and check the Routing Service.

Symptom: The Intercloud Fabric link operational status goes DOWN while the Routing Service create request is in progress.

```
Routing Service <name> failed to create since link <name> is not operational.
Failed to create the Routing Service instance <management interface IP address> for
ICF cloud.
This is primarily due to network connectivity issue between ICF and ICS.
```

Possible Cause: The operational status of the Intercloud Fabric link for the Intercloud Fabric cloud on which the Routing Service is to be created is not UP when the Routing Service create request is in progress. Once the Routing Service create request starts, Intercloud Fabric waits a maximum of 6 minutes for the Intercloud Fabric link operational status to go back UP.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Review ICFC. Verify that the 'proxyManagedEndpointStatus' MO, corresponding to the Routing Service instance, is not present in ICFC.
3. Review the ICS instance. Review the Routing Service Status and verify the associated ICS instance, when reachable, to confirm that the Routing Service container has not been created.
4. Fix the Intercloud Fabric link operational DOWN status.
5. After the Intercloud Fabric link is UP, disable the Routing Service on the Intercloud Fabric cloud.
6. Once successfully disabled, reenable the Routing Service on the Intercloud Fabric cloud.

Symptom: The Routing Service creation request from ICFC, to the associated ICS instance, fails.

```
Failed to create the Routing Service instance <management interface IP address> for
ICF cloud.
This is primarily due to network connectivity issue between ICF and ICS.
```

Possible Cause: The Intercloud Fabric link operation is UP but the Routing Service instance creation request from ICFC, to the associated ICS instance, fails due to one of the following:

- Intermittent Network connectivity issue between ICFC and that ICS instance.
- Internal error within that ICS instance because it failed to successfully process the Routing Service instance creation request from ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Review ICFC. Verify that the 'proxyManagedEndpointStatus' MO, corresponding to the Routing Service instance, is present in ICFC with the value BOOTSTRAP_FAILED.
3. Review ICS. Review the Routing Service Status and verify the associated ICS instance, when reachable, to confirm that the Routing Service container has not been created.
4. If the Intercloud Fabric link is operationally UP, confirm the operationally UP status.
5. If the Intercloud Fabric link is UP, confirm that the issue is not due to intermittent connectivity issues by attempting to ping the ICS provider public IP address.
6. If there are no intermittent Network connectivity issues, determine whether or not ICS has any internal errors when attempting to process the Routing Service create request.
7. After fixing the root cause, disable the Routing Service on the Intercloud Fabric cloud.
8. Once successfully disabled, reenable the Routing Service on the Intercloud Fabric cloud.

Symptom: ICFC fails to log in to the successfully created Routing Service instance.

```
Failed to create the Routing Service instance <management interface IP address> for
ICF cloud.
This is primarily due to network connectivity issue between ICF and ICS.
```

Possible Cause: The Routing Service instance creation request from ICFC, to the associated ICS instance, succeeds. However, ICFC is unable to successfully complete the initial handshake with that Routing Service instance due to one of the following:

- Intermittent Network connectivity issue between ICFC and that Routing Service instance.
- Internal error within the Routing Service instance because it failed to successfully process the initial handshake from ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Review ICFC. Verify that the 'proxyManagedEndpointStatus' MO, corresponding to the Routing Service instance, is present in ICFC with the value BOOTSTRAP_FAILED.
3. Review ICS. Review the Routing Service Status and verify the associated ICS instance to confirm that the Routing Service container is present.
4. If the Intercloud Fabric link is operationally UP, confirm the operationally UP status.
5. If the Intercloud Fabric link is UP, confirm that the issue is not due to intermittent connectivity issues by attempting to ping the Routing Service instance management IP address.
6. If there are no intermittent Network connectivity issues, determine whether or not the Routing Service has any internal errors when attempting to process the log in request over RESTful API.
7. After fixing the root cause, disable the Routing Service on the Intercloud Fabric cloud.
8. Once successfully disabled, reenable the Routing Service on the Intercloud Fabric cloud.

Issues Disconnecting Intercloud Fabric Routing Service from Networks

Symptom: Interface deletion errors appear when disconnecting the Routing Service from networks.

Possible Cause:

- An IP address release fails.
- Deleting an ICFC network interface fails.
- ICFC times out.
- The Routing Service is not in the system or is in a failed state.
- The Intercloud Fabric Switch or the integrated gateway cannot be reached.

Verification and Solution:

1. Review the service request logs for warnings.
2. Confirm that the Routing Service interface is properly configured and that the ICFC IP address can be reached.
3. Confirm that the Routing Service and the Intercloud Fabric Switch can be reached.
4. Confirm that the Intercloud Fabric Switch is registered with the cVSM.

Symptom: Routing Service interfaces cannot be deleted.

Possible Cause:

- All network interfaces are Layer 3 interfaces.
- The Routing Service is not in the system or is in a failed state.

Verification and Solution:

1. Review the service request logs for warnings.
2. Confirm that the data Layer 3 network (other than management and transport) is associated with the VDC.

Symptom: The Intercloud Fabric link operational status goes DOWN while the Routing Service delete request is in progress.

```
Failed to create the Routing Service instance <Routing Service IP address> for ICF
cloud. This is primarily due to network connectivity issue between ICFC and ICS.
```

Possible Cause: The operational status of the Intercloud Fabric link for the Intercloud Fabric cloud on which the Routing Service is to be deleted is not UP when the Routing Service disable request is in progress. Once the Routing Service disable request starts, Intercloud Fabric waits a maximum of 6 minutes for the Intercloud Fabric link operational status to go back UP.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Review the Intercloud Fabric workflow debug logs.
3. Review the ICFC logs.
4. Fix the Intercloud Fabric link operational DOWN status.
5. After the Intercloud Fabric link is UP, retry disabling the Routing Service on the Intercloud Fabric cloud.

Symptom: The Routing Service delete request from ICFC, to the associated ICS instance, fails.

Possible Cause: The Intercloud Fabric link operation is UP but the Routing Service instance deletion request from ICFC, to the associated ICS instance, fails due to one of the following:

- Intermittent Network connectivity issue between ICFC and that ICS instance.
- Internal error within the ICS instance due to it failing to successfully process the Routing Service instance deletion request from ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Review the Intercloud Fabric workflow debug logs.
3. Review the ICFC logs.
4. If the Intercloud Fabric link is operationally UP, confirm the operationally UP status.
5. If the Intercloud Fabric link is UP, confirm that the issue is not due to intermittent connectivity issues by attempting to ping the ICS provider public IP address.
6. If there are no intermittent Network connectivity issues, determine whether or not ICS has any internal errors when attempting to process the Routing Service delete request.
7. After fixing the root cause, disable the Routing Service on the Intercloud Fabric cloud.

Issues with Intercloud Fabric Routing Policy

Symptom: Configuring a routing policy fails when the Routing Service is being created.

Possible Cause: The Routing Service login fails.

Verification and Solution: Review the service request logs for warnings.

Symptom: Updating the routing policy is successful but the workflow fails.

```
Routing Policy Not Configured. <error string>.
```

Possible Cause:

- The workflow did not properly start.
- ICFC times out or returns an error.

Verification and Solution:

1. From the **Manage Service Request** screen, verify that the service request named **icfRoutingPolicyConfiguration** started.
2. Review the service request logs for warnings:

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending a policy create request to ICFC.

```
Routing Policy Not Configured. <error-details>.
```

Possible Cause: Intercloud Fabric fails to send the policy create request to ICFC. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, redeploy the infrastructure.

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending an update request to ICFC.

```
Routing Policy Update Failed.
```

Possible Cause: Intercloud Fabric fails to send the routing policy update request to ICFC. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, perform another routing policy update.

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending an update request to ICFC.

```
Routing Policy Update Failed.
```

Possible Cause: Intercloud Fabric fails to send the routing policy update request to ICFC. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, perform another routing policy update.

Symptom:ICFC fails to send the routing policy configuration to Routing Service instance(s).

Possible Cause: The configuration attempt fails due to one of the following reasons: Intercloud Fabric fails to send the routing policy update request to ICFC. This can be due to one of the following reasons:

1. Intermittent network connectivity issue between ICFC and that Routing Service instance.
2. An internal error within Routing Service instance occurs due to it failing to successfully process the request from ICFC.
3. The Intercloud Fabric link operational status is DOWN.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Verify if the Intercloud Fabric link is operationally UP.
3. If the Intercloud Fabric link is UP, verify that the issue is not due to intermittent connectivity issues between ICFC and the Routing Service.
4. If the Intercloud Fabric link is operationally UP and there are no intermittent network connectivity issues, check if the Routing Service has any internal errors when attempting to process the request.
5. Once the root cause is fixed, ICFC automatically pushes the latest routing policy configuration to the Routing Service instance(s).
6. Intercloud Fabric detects the successful configuration of the routing policy by verifying it with ICFC. It then updates the routing policy configuration status for the Routing Service instance to be configured.

Symptom: Intercloud Fabric and ICFC suffer a communication error while polling to verify the results of policy configuration on the Routing Service instance(s).

```
Routing Service Update Failed [ <link name>: ErrorCode:<code> Details:<ICFC
communication exception message> ].
```

Possible Cause: Intercloud Fabric fails to receive the Routing Service instance status to verify if the configuration attempt is successful or not. This can be due to the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, ICFC automatically pushes the latest routing policy configuration to the Routing Service instance(s).
4. Intercloud Fabric detects the successful configuration of the routing policy by verifying it with ICFC. It then updates the routing policy configuration status for the Routing Service instance to be configured.

Symptom: Intercloud Fabric times out waiting for ICFC to configure the policy on its Routing Service instance(s).

```
Update Routing Policy FailedConfiguring Routing Policy on Routing Service on ICF link
 <link name> timed out: <poll timeout message in ICF>.
```

Possible Cause: Intercloud Fabric gives up waiting for ICFC to configure the routing policy. ICFC suffers an internal error. This prevents ICFC itself from timing out the policy configuration flow.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to ICFC failing to complete the operation in time by restarting ICFC services.
3. Once the root cause is fixed, ICFC automatically pushes the latest routing policy configuration to the Routing Service instance(s).
4. Intercloud Fabric detects the successful configuration of the routing policy by verifying it with ICFC. It then updates the routing policy configuration status for the Routing Service instance to be configured.

    - Network delete request.
    - Network update request (dummy update).
    - VDC delete request.

Issues with Intercloud Fabric Virtual Data Center

Symptom: Traffic between cloud VMs fails.

Possible Cause: The Routing Service fails to connect the cloud VMs.

Verification and Solution:

1 Confirm that the Routing Service can reach both of the cloud VM NICs by pinging the cloud VM IP addresses from the Routing Service.

Perform a test ping:

```
<Menu Options list>

SELECT> 1

Enter IP Adress/Hostname:
```

After the IP address is entered, the **ping** command is executed for 5 packets:

```
Enter IP Address/Hostname : 10.36.10.198
PING 10.36.10.198 (10.36.10.198) 56(84) byts of data.
64 bytes from 10.36.10.198: icmp_seq=1 ttl=64 time=5.28 ms
64 bytes from 10.36.10.198: icmp_seq=2 ttl=64 time=5.78 ms
64 bytes from 10.36.10.198: icmp_seq=3 ttl=64 time=5.38 ms
64 bytes from 10.36.10.198: icmp_seq=4 ttl=64 time=6.66 ms
64 bytes from 10.36.10.198: icmp_seq=5 ttl=64 time=5.35 ms

--- 10.36.10.198 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 5.285/5.694/6.660/0.518 ms

Press return to continue ...
```

2 If the connection to the Routing Service fails, determine whether or not the site-to-site tunnel status on the Intercloud Fabric Switch is operational.

```
AWS-712-Link-ics-1# show intercloud tunnel brief
--------------------------------------------------------------
Tunnel   Tunnel          Tunnel   Control   Peer
ID       Type            Status   Status    Name
--------------------------------------------------------------
1        Site-to-Site    Up       Up        AWS-712-Link-ics-1
```

3 Confirm that the Intercloud Fabric Switch is registered with the cVSM.

```
vsm598859926# show mod
Mod Ports Module-Type                   Model              Status
--- ----- ------------------------      -----------        ----------
1   0     Virtual Supervisor Module     Nexus1000V         active*
3   1002  Virtual Cloud Extender Module InterCloudExtender ok
4   1022  Virtual Cloud Extender Module InterCloudSwitch   ok

Mod Sw              Hw
--- -----------     -----------------------
1   5.2(1)SK3(1.3)  0.0
3   5.2(1)SK3(1.3)  Linux 3.14.27-ics10
4   5.2(1)SK3(1.3)  Linux 3.14.27-ics10

Mod Server-IP         Server-UUID                           Server-Name
--- -------------     ------------------------              ----------------
1   10.36.5.90        NA                                    NA
3   10.36.5.91        2F50E5A0-47DF-5877-3C6D-3573F3351C94 AWS-712-Link-icx-1
4   10.36.5.92        485642DC-00AB-C149-F5CF-088930A4ADBB AWS-712-Link-icx-1

* this terminal session
```

4 Determine whether or not the secure tunnel to cloud VM is operational on the Intercloud Fabric Switch.

```
AWS-712-Link-ics-1# show intercloud tunnel brief
--------------------------------------------------------------------
Tunnel    Tunnel          Tunnel    Control    Peer
ID        Type            Status    Status     Name
--------------------------------------------------------------------
```

```
1         Site-to-Site    Up        Up         AWS-712-Link-ics-1
2         Access          Up        Up         vm-a27f5
3         Access          Up        Up         vm-6e5de
4         Access          Up        Up         vm-17029
5         Access          Up        Up         vm-24875
6         Access          Up        Up         vm-ae865
--------------------------------------------------------------------
```

5 In the Routing Service, verify that the interfaces are created for both networks and are in an operational state.

6 Review the display network details to confirm that the data VLAN for both VLANs is in operation and contains the IP addresses from their respective networks.

```
Interfaces
-----------
data
ip6gre0
ip6tn10
lo
mgmt
pCeth

Enter the interfaces name details. Press return for all interfaces, or 0 to exit.
```

Output for details on 1 interface:

```
Enter the interfaces name details. Press return for all interfaces, or 0 to exit:
mgmt

Interface details
-----------------
13: mgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> qdisc pfifo_fast state UP mode DEFAULT
 group default qlen 1000
    link/either 7e:00:18:70:00:03 brd ff:ff:ff:ff:ff:ff promiscuity 0
    veth
    RX: bytes  packets  errors  dropped overun  mcast
    2243384597 13757572 0       4870857 0       0
    TX: bytes  packets  errors  dropped carrier collsns
    4676361    29527    0       0       0       0

Address
-------
13: mgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> qdisc pfifo_fast state UP group default
 qlen 1000
    link/either 7e:00:18:70:00:03 brd ff:ff:ff:ff:ff:ff
    inet 10.36.70.28/16 brd 10.36.255.255 scope global mgmt
      valid_lft forever preferred_lft forever
    inet6 fe80::7c00::18ff:fe70:3/64 scope link
      valid_lft forever preferred_lft forever
```

Output for all interfaces:

```
data    Link encap:Ethernet  HWaddr 7e:00:18:70:00:04
        inet6 addr: fe80::7c00:18ff:fe70:4/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1352 Metric:1
        RX packets:13724218 errors:0 dropped:4871814 overruns:0 frame:0
        TX packets:9 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:2239552998 (2.2 GB)  TX bytes:738 (738.0 B)

ip6gre0 Link ecap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        NOARP  MTU:1448  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 GB)  TX bytes:0 (0.0 B)

ip6tn10 Link ecap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        NOARP  MTU:1452  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 GB)  TX bytes:0 (0.0 B)

lo      Link ecap:Local loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        inet6 addr: ::1/128  Scope:Host
        UP LOOPBACK RUNNING  MTU:65536 Metric:1
        RX packets:3473237 errors:0 dropped:0 overruns:0 frame:0
        TX packets:3473237 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:181018327 (181.0 GB)  TX bytes:181018327 (181.0 B)

mgmt    Link encap:Ethernet  HWaddr 7e:00:18:70:00:03
        inet addr:10:36:70:28  Bcast:10:36:255:255  Mask:255:255:0:0
        inet6 addr: fe80::7c00:18ff:fe70:3/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1352 Metric:1
        RX packets:13759826 errors:0 dropped:4871612 overruns:0 frame:0
        collisions:0 txqueuelen:1000
        RX bytes:2243746072 (2.2 GB)  TX bytes:4678855 (4.6 B)

pCeth   Link encap:Ethernet  HWaddr a2:e6:22:1a:20:f2
        inet addr:172:17:10:2  Bcast:172:17:10:255  Mask:255:255:0
        BROADCAST MULTICAST MTU:1500 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 GB)  TX bytes:0 (0.0 B)
```

7 Log in to the Intercloud Fabric Switch and verify that the data trunk port-profile of 0-data-vem has the VLANs of both cloud VM NICs.

```
link1-ics-1# show intercloud services
Routing
===========================================================
Enabled: Yes
Status: Running <--- Status should be Running
Management IP:10:193:180:237

link1-ics-1# vemcmd show port
  LTL  VSM Port Admin Link State PC-LTL SGID   Vem Port Type ORG svcpath Owner
   50     Veth3    UP   UP   FWD      0        eth1-vem        1       0
   51     Veth4    UP   UP   F/B*     0         veth1-0        0       0 <---+
   52     Veth5    UP   UP   F/B*     0      0-data-vem        0       0 <---+
Should be in F/B state. Note VSM Port name for the next command
   53     Veth6    UP   UP   FWD      0      0-mgmt-vem        1       0

F/B: Port is BLOCKED on some of the vlans.
      One or more vlans are either not created or
      not in the list of allowed vlans for this port.
 Please run "vemcmd show port vlans" to see the details.
link1-ics-1# vemcmd show port vlans
                         Native   VLAN   Allowed
  LTL  VSM Port  Mode     VLAN     State* Vlans
   50     Veth3   A        280    FWD    280
   51     Veth4   T          1    FWD    280,
   52     Veth5   T          1    FWD    280, 1161-1163 <--- Should have the VLANs
 of both the CVMs.
   53     Veth6   A        280    FWD    280

VLAN State: VLAN State represents the state of allowed vlans.
link1-ics-1# vemcmd show dvport
  LTL  VSM Port  DVPortID      DVPortGroup  Vem Port
   50     Veth3         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Access_36
eth1-vem
   51     Veth4         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Trunk_Tunnel
veth1-0
   52     Veth5         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Trunk_Tunnel
0-data-vem <-- Check this port profile in VSM has all the L3 vlan-s of CVSM.
   53     Veth6         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Access_36
0-mgmt-vem
```

8 If 0-data-vem is in BLK state, determine whether or not the Intercloud Fabric Switch is registered with the cVSM.

9 If 0-data-vem does not include both VLANs, determine whether or not the VLANs have been created in the cVSM and are part of the trunk port-profile.

10 If the cloud VM network is stretched across Layer 3, verify that the ARP filter entry is created in the Intercloud Fabric Switch for the private cloud gateway with the MAC address of the Routing Service interface.

```
link1-ics-1#vemcmd show arp all
Flags: D-Dynamic S-Static d-Delete s-Sticky P-Proxy B-Public C-Create
X-Exlusive
VLAN/SEFID     IP Address     MAC Address        Flags     Expiry
1161          192.168.61.1   8e.00:04:00:00:04  S P       0   <--- Make sure that
 the flags S and P are set.
```

11 Log in to the cloud VM and verify that it is configured with the correct IP address for the default gateway (**ip route show** on Linux cloud VMs, **route print** on Windows cloud VMs). For Layer 3 cloud-only networks, the default gateway should be the same as the cloud gateway; for networks stretched across Layer 3, the default gateway should be the private cloud gateway.

12 Log in to the cloud VMs and verify that they are configured with the private cloud gateway IP address for the default gateway (**ip route show** on Linux cloud VMs, **route print** on Windows cloud VMs). For Layer 3 cloud-only networks, the default gateway should be the same as the cloud gateway; for networks stretched across Layer 3, the default gateway should be the private cloud gateway.

Symptom: Traffic between the cloud VM and the private cloud VM fails.

Possible Cause: The Routing Service fails to connect the cloud VM and the private cloud VM IP addresses from the Routing Service.

Verification and Solution:

1. Confirm that the Routing Service can reach both of the cloud VM NICs by pinging the cloud VM IP addresses from the Routing Service.

   Perform a test ping:

   ```
   <Menu Options list>

   SELECT> 1

   Enter IP Adress/Hostname:
   ```

   After the IP address is entered, the **ping** command is executed for 5 packets:

   ```
   Enter IP Address/Hostname : 10.36.10.198
   PING 10.36.10.198 (10.36.10.198) 56(84) byts of data.
   64 bytes from 10.36.10.198: icmp_seq=1 ttl=64 time=5.28 ms
   64 bytes from 10.36.10.198: icmp_seq=2 ttl=64 time=5.78 ms
   64 bytes from 10.36.10.198: icmp_seq=3 ttl=64 time=5.38 ms
   64 bytes from 10.36.10.198: icmp_seq=4 ttl=64 time=6.66 ms
   64 bytes from 10.36.10.198: icmp_seq=5 ttl=64 time=5.35 ms

   --- 10.36.10.198 ping statistics ---
   5 packets transmitted, 5 received, 0% packet loss, time 4006ms
   rtt min/avg/max/mdev = 5.285/5.694/6.660/0.518 ms

   Press return to continue ...
   ```

2. Check if the site-to-site tunnel status on the Intercloud Fabric Switch is operational.

   ```
   AWS-712-Link-ics-1# show intercloud tunnel brief
   --------------------------------------------------------------
   Tunnel   Tunnel          Tunnel   Control   Peer
   ID       Type            Status   Status    Name
   --------------------------------------------------------------
   1        Site-to-Site    Up       Up        AWS-712-Link-ics-1
   ```

3. Check if the Intercloud Fabric Switch is registered with the cVSM.

   ```
   vsm598859926# show mod
   Mod Ports Module-Type                   Model              Status
   --- ----- ------------------------      -----------        ----------
   1   0     Virtual Supervisor Module     Nexus1000V         active*
   3   1002  Virtual Cloud Extender Module InterCloudExtender ok
   4   1022  Virtual Cloud Extender Module InterCloudSwitch   ok

   Mod Sw               Hw
   --- -----------      -----------------------
   1   5.2(1)SK3(1.3)   0.0
   3   5.2(1)SK3(1.3)   Linux 3.14.27-ics10
   4   5.2(1)SK3(1.3)   Linux 3.14.27-ics10

   Mod Server-IP         Server-UUID                          Server-Name
   --- -------------     ------------------------             ----------------
   1   10.36.5.90        NA                                   NA
   3   10.36.5.91        2F50E5A0-47DF-5877-3C6D-3573F3351C94 AWS-712-Link-icx-1
   4   10.36.5.92        485642DC-00AB-C149-F5CF-088930A4ADBB AWS-712-Link-icx-1

   * this terminal session
   ```

4. Confirm that the secure tunnel to the cloud VM is operational on the Intercloud Fabric Switch.

   ```
   AWS-712-Link-ics-1# show intercloud tunnel brief
   --------------------------------------------------------------------
   Tunnel    Tunnel          Tunnel    Control    Peer
   ID        Type            Status    Status     Name
   --------------------------------------------------------------------
   1         Site-to-Site    Up        Up         AWS-712-Link-ics-1
   2         Access          Up        Up         vm-a27f5
   3         Access          Up        Up         vm-6e5de
   ```

```
4         Access          Up        Up          vm-17029
5         Access          Up        Up          vm-24875
6         Access          Up        Up          vm-ae865
-----------------------------------------------------------------
```

5 If the private cloud VM is on a stretched network, verify in the Routing Service that interfaces are created for both networks and are in an operational state.

6 Review the display network details to confirm that the data VLAN for both VLANs is in operation and contains the IP addresses from their respective networks.

```
Interfaces
-----------
data
ip6gre0
ip6tn10
lo
mgmt
pCeth

Enter the interfaces name details. Press return for all interfaces, or 0 to exit.
```

Output for details on 1 interface:

```
Enter the interfaces name details. Press return for all interfaces, or 0 to exit:
mgmt

Interface details
-----------------
13: mgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> qdisc pfifo_fast state UP mode DEFAULT
 group default qlen 1000
    link/either 7e:00:18:70:00:03 brd ff:ff:ff:ff:ff:ff promiscuity 0
    veth
    RX: bytes  packets  errors  dropped overun  mcast
    2243384597 13757572 0       4870857 0       0
    TX: bytes  packets  errors  dropped carrier collsns
    4676361    29527    0       0       0       0

Address
-------
13: mgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> qdisc pfifo_fast state UP group default
 qlen 1000
    link/either 7e:00:18:70:00:03 brd ff:ff:ff:ff:ff:ff
    inet 10.36.70.28/16 brd 10.36.255.255 scope global mgmt
      valid_lft forever preferred_lft forever
    inet6 fe80::7c00::18ff:fe70:3/64 scope link
      valid_lft forever preferred_lft forever
```

Output for all interfaces:

```
data    Link encap:Ethernet  HWaddr 7e:00:18:70:00:04
        inet6 addr: fe80::7c00:18ff:fe70:4/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1352 Metric:1
        RX packets:13724218 errors:0 dropped:4871814 overruns:0 frame:0
        TX packets:9 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:2239552998 (2.2 GB)  TX bytes:738 (738.0 B)

ip6gre0 Link ecap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        NOARP  MTU:1448  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 GB)  TX bytes:0 (0.0 B)

ip6tn10 Link ecap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        NOARP  MTU:1452  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 GB)  TX bytes:0 (0.0 B)

lo      Link ecap:Local loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        inet6 addr: ::1/128  Scope:Host
        UP LOOPBACK RUNNING  MTU:65536 Metric:1
        RX packets:3473237 errors:0 dropped:0 overruns:0 frame:0
        TX packets:3473237 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:181018327 (181.0 GB)  TX bytes:181018327 (181.0 B)

mgmt    Link encap:Ethernet  HWaddr 7e:00:18:70:00:03
        inet addr:10:36:70:28  Bcast:10:36:255:255  Mask:255:255:0:0
        inet6 addr: fe80::7c00:18ff:fe70:3/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1352 Metric:1
        RX packets:13759826 errors:0 dropped:4871612 overruns:0 frame:0
        collisions:0 txqueuelen:1000
        RX bytes:2243746072 (2.2 GB)  TX bytes:4678855 (4.6 B)

pCeth   Link encap:Ethernet  HWaddr a2:e6:22:1a:20:f2
        inet addr:172:17:10:2  Bcast:172:17:10:255  Mask:255:255:0
        BROADCAST MULTICAST MTU:1500 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 GB)  TX bytes:0 (0.0 B)
```

7. If the private cloud VM is on a private cloud-only network, confirm that the transport network interface is created on the Routing Service and is in an operational state.

8. Review the display routing table details to confirm that a route exists to the private cloud VM network by way of the transport network gateway.

9. Log in to the Intercloud Fabric Switch and verify that the data trunk port-profile of 0-data-vem includes the VLAN of the cloud VM.

10. If the private cloud VM is on a stretched network, verify that the data trunk port-profile that contains 0-data-vem and veth 1-0 includes the VLAN of the private cloud VM.

11 If the private cloud VM is on a private cloud-only network, verify that 0-data-vem and veth 1-0 include the VLAN of the transport network.

```
link1-ics-1# show intercloud services
Routing
==============================================================
Enabled: Yes
Status: Running <--- Status should be Running
Management IP:10:193:180:237

link1-ics-1# vemcmd show port
  LTL  VSM Port Admin Link State PC-LTL SGID   Vem Port Type ORG svcpath Owner
   50     Veth3    UP   UP   FWD      0        eth1-vem        1       0
   51     Veth4    UP   UP   F/B*     0         veth1-0        0       0 <---+
   52     Veth5    UP   UP   F/B*     0      0-data-vem        0       0 <---+
Should be in F/B state. Note VSM Port name for the next command
   53     Veth6    UP   UP   FWD      0      0-mgmt-vem        1       0

F/B: Port is BLOCKED on some of the vlans.
      One or more vlans are either not created or
      not in the list of allowed vlans for this port.
 Please run "vemcmd show port vlans" to see the details.
link1-ics-1# vemcmd show port vlans
                         Native   VLAN   Allowed
  LTL  VSM Port  Mode    VLAN     State* Vlans
   50     Veth3   A        280    FWD    280
   51     Veth4   T          1    FWD    280, 1162 <--- Should have the VLAN
EVM/transport network
   52     Veth5   T          1    FWD    280, 1161-1163 <--- Should have the VLANs
 of both the CVM and EVM/transport network.
   53     Veth6   A        280    FWD    280

VLAN State: VLAN State represents the state of allowed vlans.
link1-ics-1# vemcmd show dvport
  LTL  VSM Port  DVPortID      DVPortGroup  Vem Port
   50     Veth3         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Access_36
eth1-vem
   51     Veth4         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Trunk_Tunnel
veth1-0
   52     Veth5         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Trunk_Tunnel
0-data-vem <-- Check this port profile in VSM has all the L3 vlan-s of both the CVM
 and EVM/transport network.
   53     Veth6         0 e52e4c4c-38e7-4774-84b0-16dc6a5063a3_ICS_Access_36
0-mgmt-vem
```

12 If 0-data-vem is in BLK state, confirm that the Intercloud Fabric Switch is registered with cVSM.

13 If 0-data-vem does not have both VLANs, determine whether or not the VLANs have been created in cVSM and are part of the trunk port-profile.

14 If the network is stretched across Layer 3, verify that the ARP filter entry is created in the Intercloud Fabric Switch for the private cloud gateway with the MAC address of the Routing Service interface. Make sure flags S and P are set.

```
link1-ics-1# vemcmd show arp all
Flags: D-Dynamic S-Static d-Delete s-Sticky P-Proxy B-Public C-Create
X-Exlusive
VLAN/SEFID     IP Address     MAC Address        Flags     Expiry
1161           192.168.61.1   8e.00:04:00:00:04  S P       0
```

15 Log in to ICX and verify that both trunk ports carry the private cloud VM VLAN.

```
link-icx-1# vemcmd show port vlans
                        NATIVE    VLAN     Allowed
LTL   VSM Port   Mode   VLAN      State*   Vlans
 50     Veth1    T      1         FWD      280, 1162
 51     Veth2    T      1         FWD      280, 1162
```

16 Log in to the cloud VM and verify that it is configured with the correct IP address for the default gateway (**ip route show** on Linux cloud VMs, **route print** on Windows cloud VMs). For Layer 3 cloud-only networks, the default gateway should be the same as the cloud gateway; for networks stretched across Layer 3, the default gateway should be the private cloud gateway.

17 Log in to the cloud VMs and verify that they are configured with the private cloud gateway IP address for the default gateway (**ip route show** on Linux cloud VMs, **route print** on Windows cloud VMs). If the private cloud VM is on private cloud-only network, verify that the transport network gateway can be reached from the private cloud VM.

Symptom: Connectivity between a private VM and a cloud VM in different VLANs fails intermittently.

Possible Causes:

- A Cisco Nexus 1000V is used in the private data center with Unknown-Unicast-Flooding-Block enabled.
- The private VM and cloud VM are in different VLANs, thereby causing inter-VLAN routing to fail.

Verification and Solution:

1 Confirm that packets are being dropped by entering the following command on the VEM:

```
vemcmd show port-drops
```

2 Determine whether Unknown-Unicast-Flooding-Block is enabled or disabled by entering the following command on the Cisco Nexus 1000V:

```
privateVSM# show run | include uufb
uufb enable
```

3 If Unknown-Unicast-Flooding-Block is enabled, disable it by entering the following command in config mode:

```
no uufb enable
```

If Unknown-Unicast-Flooding-Block is already disabled, see other topics in this section for more information about troubleshooting inter-VLAN routing issues.

Symptom: For cloud VMs whose default gateway points to the Routing Service that is in the cloud, inter-VLAN routing does not work, but communication within the same VLAN does work.

Possible Cause:

- The Routing Service subinterface IP addresses are incorrect.
- The default gateway for the cloud VM does not point to the Routing Service subinterface IP address.
- The Routing Service does not have the appropriate routes.
- The Intercloud Fabric Switch is routing a VLAN that is not in the list of allowed VLANs in the Cisco Nexus 1000V default trunk.
- The cloud VM cannot reach the default gateway.

Verification and Solution:

1. In ICFC, ensure that the Routing Service subinterfaces are configured with the correct IP address.
2. Log in to the cloud VM and verify that its default gateway points to the Routing Service subinterface IP address (from Step 1).
3. Log in to the Routing Service and verify that the appropriate routes are present.
4. Log in to the Intercloud Fabric Switch and verify that the VLAN being routed is in the list of allowed VLANs in the Cisco Nexus 1000V default trunk:

```
# vemcmd show port
# vemcmd show port vlans
```

5. Verify that the cloud VM can reach the default gateway, which is the Routing Service subinterface.

Symptom: For cloud VMs whose default gateway points to the Routing Service that is in the private data center, inter-VLAN routing does not work, but communication within the same VLAN does work.

Possible Cause:

- The Routing Service subinterface IP addresses are incorrect.
- The Intercloud Fabric Switch is routing a VLAN that is not in the list of allowed VLANs in the Cisco Nexus 1000V default trunk.
- The Routing Service does not have the appropriate routes.
- The default gateway for the cloud VM is a Routing Service that is not in the cloud.
- The cloud VM cannot reach the default gateway.
- Inter-VLAN routing is not working.

Verification and Solution:

1. In ICFC, confirm that the **Extend Default Gateway** option is checked for the Routing Service subinterface and that the IP addresses are correct.
2. Log in to the Intercloud Fabric Switch and verify that the VLAN being routed is in the list of allowed VLANs in the Cisco Nexus 1000V default trunk:

```
# vemcmd show port
# vemcmd show port vlans
```

3. From the Routing Service, confirm that the IP address from Step 1 is displayed in the output of the following command:

```
# vemcmd show arp all
```

4. Log in to the Routing Service and verify that the appropriate routes are present.
5. In ICFC, uncheck the **Extend Default Gateway** box.
6. Log in to the cloud VM and change the default gateway to the Routing Service that is in the cloud.
7. Verify that the cloud VM can reach the default gateway, which is the Routing Service subinterface.
8. Verify that inter-VLAN routing works.

Issues with Intercloud Fabric Routing Service Internal Errors

Symptom: Intercloud Fabric fails to communicate with ICFC.

```
Routing Service creation request to ICFC failed. <ICFC exception message>.
```

Possible Cause: Intercloud Fabric fails to successfully communicate with ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Review the Intercloud Fabric workflow debug logs.
3. Fix the root cause failure and disable the Routing Service on the Intercloud Fabric cloud.
4. Once successfully disabled, reenable the Routing Service on that Intercloud Fabric cloud.

Symptom: Intercloud Fabric times out during the Routing Service create instance.

```
Routing Service creation has timed out. ICFC internal error.
```

Possible Cause: Intercloud Fabric times out waiting for ICFC to create the Routing Service instance due to an ICFC internal error that prevents ICFC to itself time out the Routing Service create instance.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the root cause failure by restarting the ICFC services.
3. Disable the Routing Service on the Intercloud Fabric cloud.
4. Once successfully disabled, reenable the Routing Service on that Intercloud Fabric cloud.

Symptom: ICFC rejects the Routing Service create request.

Possible Cause: Intercloud Fabric and ICFC are not synchronized after a Routing Service is deleted without being successfully cleaned up from ICFC and a new Routing Service is created with the same name (same as the Intercloud Fabric link name).

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Do one of the following:
    - Create the Intercloud Fabric cloud and Intercloud Fabric link with different names.
    - Contact Cisco Customer Support to assist in getting the ICFC database to be synchronized with Intercloud Fabric.

Symptom: Intercloud Fabric times out during the Routing Service delete instance.

Possible Cause: Intercloud Fabric times out waiting for ICFC to delete the Routing Service instance due to an ICFC internal error that prevents ICFC to itself time out the Routing Service delete instance.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Contact Cisco Customer Support to debug and fix the ICFC internal error problem.
3. After fixing the root cause, perform one of the following:
    - If the Intercloud Fabric link deletion request failed, delete the Intercloud Fabric link again.
    - If the Routing Service disable request failed, disable it again.

Issues with the Intercloud Fabric Routing Service Interface

This section includes symptoms, possible causes, and solutions for issues associated with the Routing Service interface.

Issues Enabling and Configuring Intercloud Fabric Routing Service Interface

Symptom: The Layer 3 data network, which is a stretched network, does not have a free IP address.

```
Invalid Interface IP for Network <network name>.
```

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Add at least one more IP pool to the network.
3. Perform a dummy update on the network without changing any of its properties.

Symptom: The Layer 3 data network, which is a cloud-only network, does not have a reserved cloud gateway IP address.

```
Invalid Interface IP for Network <network name>.
```

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Delete and recreate the cloud-only network.

Issues Disabling Intercloud Fabric Routing Service Interface

Symptom: The Routing Service interface IP address release fails.

```
Failed to release Mgmt IP <ips> for IG Intance <Routing Service name> (error message).
```

Possible Cause: The allocation failure can occur on one or more of the following interface types:

- Management network interface.
- Transport network interface.
- Data-only network interface.

Verification and Solution: Review the Intercloud Fabric service request logs.

Issues Connecting Intercloud Fabric Routing Service Interface to Networks

Symptom: ICFC fails to send the interface create request to the Routing Service instance(s).

```
Failed to configure Routing Service interface for Network (<network name>) on ICFCloud
 (<cloud name>). Failed to configure Routing Service interface.
```

Possible Cause: The interface create request, from Intercloud Fabric to ICFC, fails even though the Intercloud Fabric link is operationally UP. The creation fails for one of the following reasons:

- Intermittent network connectivity between ICFC and that Routing Service instance.
- An internal error within the Routing Service instance fails to successfully process the request from ICFC.
- Prior to the Intercloud Fabric link being sent to the Routing Service instance(s), its operational status goes DOWN after the interface create request is sent to ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Verify that the Intercloud Fabric link is operationally UP.
3. If the Intercloud Fabric link is UP, make sure that the issue is not due to intermittent connectivity issues between ICFC and the Routing Service.
4. If the Intercloud Fabric link is operationally UP and there are no intermittent network connectivity issues, determine if the Routing Service has any internal errors when attempting to process the request.
5. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:

    - Network update request (dummy update): network create/update as the initial request.
    - VDC delete followed by VDC create: VDC create as the initial request.

Symptom: ICFC fails to successfully send the ARP filter create request to the Intercloud Fabric Switch instance(s).

```
Failed to configure Routing Service interface for Network (<network name>) on ICFCloud
 (<cloud name>). Enabling proxy ARP for Gateway (<enterprise gateway IP>) in ICS failed.
```

Possible Cause: The ARP filter create request, from Intercloud Fabric to ICFC, fails even though the Intercloud Fabric link is operationally UP. the creation fails for one of the following reasons:

- Intermittent network connectivity between ICFC and that Intercloud Fabric Switch instance.
- An internal error within the Intercloud Fabric Switch instance fails to successfully process the request from ICFC.
- Prior to the Intercloud Fabric link being sent to the Intercloud Fabric Switch instance(s), its operational status goes DOWN after the interface create request is sent to ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Verify that the Intercloud Fabric link is operationally UP.
3. If the Intercloud Fabric link is UP, make sure that the issue is not due to intermittent connectivity issues between ICFC and the Intercloud Fabric Switch.
4. If the Intercloud Fabric link is operationally UP and there are no intermittent network connectivity issues, determine if Intercloud Fabric Switch has any internal errors when attempting to process the request.
5. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:

    - Network update request (dummy update): network create/update as the initial request.
    - VDC delete followed by VDC create: VDC create as the initial request.

Issues Disconnecting Intercloud Fabric Routing Service Interface from Networks

Symptom: ICFC fails to successfully send the ARP filter delete request to the Intercloud Fabric Switch instance(s).

```
Disabling proxy ARP in ICS for Gateway (<enterprise gateway IP>) failed.
```

Possible Cause: The ARP filter delete request, from Intercloud Fabric to ICFC, fails even though the Intercloud Fabric link is operationally UP. The deletion fails for one of the following reasons:

- Intermittent network connectivity between ICFC and that Intercloud Fabric Switch instance.
- An internal error within the Intercloud Fabric Switch instance fails to successfully process the request from ICFC.
- Prior to the Intercloud Fabric link being sent to the Intercloud Fabric Switch instance(s), its operational status goes DOWN after the ARP filter delete request is sent to ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Verify that the Intercloud Fabric link is operationally UP.
3. If the Intercloud Fabric link is UP, make sure that the issue is not due to intermittent connectivity issues between ICFC and the Intercloud Fabric Switch.
4. If the Intercloud Fabric link is operationally UP and there are no intermittent network connectivity issues, determine if Intercloud Fabric Switch has any internal errors when attempting to process the request.
5. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:

   - Network delete request.
   - Network update request (dummy update).
   - VDC delete request.

Symptom: ICFC fails to successfully send the interface delete request to the Routing Service instance(s).

```
Deleting Routing Service interface from Routing Service for Network (<network name>)
failed. Failed to delete Routing Service interface.
```

```
Deleting Routing Service interface from Routing Service for Network (<network name>)
failed. Failed to delete Routing Service interface.
```

Possible Cause: The interface delete request, from Intercloud Fabric to ICFC, fails even though the Intercloud Fabric link is operationally UP. The deletion fails for one of the following reasons:

- Intermittent network connectivity between ICFC and that Routing Service instance.
- An internal error within the Routing Service instance fails to successfully process the request from ICFC.
- Prior to the Intercloud Fabric link being sent to the Routing Service instance(s), its operational status goes DOWN after the interface create request is sent to ICFC.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Verify that the Intercloud Fabric link is operationally UP.
3. If the Intercloud Fabric link is UP, make sure that the issue is not due to intermittent connectivity issues between ICFC and the Routing Service.
4. If the Intercloud Fabric link is operationally UP and there are no intermittent network connectivity issues, determine if the Routing Service has any internal errors when attempting to process the request.
5. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:
    - Network delete request.
    - Network update request (dummy update).
    - VDC delete request.

Issues with Intercloud Fabric Routing Service Interface Internal Creation Errors

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending an interface create request to ICFC.

```
Failed to create Routing Service interface for Network (<network name) on
ICFCloud(<cloud name>). Creating Routing Service Interface for Network (network name>)
 failed. Error: <message from ICFC>.
```

Possible Cause: Intercloud Fabric fails to send the interface create request to ICFC. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:

- Network update request (dummy update): network create/update as initial request.
- VDC delete followed by VDC create: VDC create an initial request.

Symptom: Intercloud Fabric and ICFC suffer a communication error when polling to review the outcome of the interface create operation in ICFC.

```
Failed to create Routing Service interface for Network (<network name>) on
ICFCloude(<cloud name>). Creating Routing Service Interface for Network (<network
name>) failed. Error: <message from ICFC>.
```

Possible Cause: Intercloud Fabric fails to receive the status of the Routing Service instance to verify if the create operation is successful or not. This is due to the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:
    - Network update request (dummy update): network create/update as initial request.
    - VDC delete followed by VDC create: VDC create an initial request.

Symptom: Intercloud Fabric times out waiting for ICFC to create the interface.

```
Failed to create Routing Service interface for Network (<network name>) on
ICFCloude(<cloud name>). Creating Routing Service Interface for Network (<network
name>) failed. Error: <message from ICFC>.
```

Possible Cause: Intercloud Fabric gives up waiting for ICFC to create the interface. This only occurs when ICFC suffers an internal error. This internal error prevents ICFC itself from timing out the interface deletion flow.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to ICFC failing to complete the operation in time by restarting ICFC services.
3. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:
    - Network update request (dummy update): network create/update as initial request.
    - VDC delete followed by VDC create: VDC create an initial request.

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending an ARP filter create request to ICFC.

```
Failed to create Routing Service interface for Network (<network name>) on
ICFCloude(<cloud name>). Creating Routing Service Interface for Network (<network
name>) failed. Error: <message from ICFC>.
```

Possible Cause: Intercloud Fabric fails to send the ARP filter create request to ICFC. This is due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:
    - Network update request (dummy update): network create/update as initial request.
    - VDC delete followed by VDC create: VDC create an initial request.

Symptom: Intercloud Fabric and ICFC suffer a communication error while polling to review the outcome of an ARP filter create operation in ICFC.

```
Failed to create Routing Service interface for Network (<network name>) on
ICFCloude(<cloud name>). Creating Routing Service Interface for Network (<network
name>) failed. Error: <message from ICFC>.
```

Possible Cause: Intercloud Fabric fails to receive the status of the Intercloud Fabric Switch instance to verify if the create operation is successful or not. This is due to the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:
    - Network update request (dummy update): network create/update as initial request.
    - VDC delete followed by VDC create: VDC create an initial request.

Symptom: Intercloud Fabric times out waiting for ICFC to create the ARP filter.

```
Failed to create Routing Service interface for Network (<network name>) on
ICFCloude(<cloud name>). Creating Routing Service Interface for Network (<network
name>) failed. Error: <message from ICFC>.
```

Possible Cause: Intercloud Fabric gives up waiting for ICFC to create the ARP filter. This only occurs when ICFC suffers an internal error. This internal error prevents ICFC itself from timing out the ARP filter deletion flow.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to ICFC failing to complete the operation in time by restarting ICFC services.
3. Once the root cause is fixed, perform the appropriate operation corresponding to the previous operation that initially triggered the interface creation. The following are the possible operations:
    - Network update request (dummy update): network create/update as initial request.
    - VDC delete followed by VDC create: VDC create an initial request.

Issues with Intercloud Fabric Routing Service Interface Internal Deletion Errors

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending an interface delete request to ICFC.

```
Deleting Routing Service interface from Routing Service for Network (<network name)
failed. Error: <message indicating type of ICFC communication exception>.
```

Possible Cause: Intercloud Fabric fails to send the interface delete request to ICFC. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:
    - Network delete request.
    - Network update request (dummy update).
    - VDC delete request.

Symptom: Intercloud Fabric and ICFC suffer a communication error while polling to review the outcome of the interface delete operation in ICFC.

```
Deleting Routing Service interface from Routing Service for Network (<network name)
failed. Error: <message indicating type of ICFC communication exception>.
```

Possible Cause: Intercloud Fabric fails to receive the status of the Routing Service instance to verify if the delete operation is successful or not. This is due to the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:

    - Network delete request.
    - Network update request (dummy update).
    - VDC delete request.

Symptom: Intercloud Fabric times out waiting for ICFC to delete the interface.

```
Deleting Routing Service interface from Routing Service for Network (<network name)
failed. Error: <timeout message generated within Intercloud Fabric>.
```

Possible Cause: Intercloud Fabricgives up waiting for ICFC to delete the interface. ICFC has suffered an internal error. This internal error prevents ICFC itself from timing out the interface deletion flow.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to ICFC failing to complete the operation in time by restarting ICFC services.
3. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:

    - Network delete request.
    - Network update request (dummy update).
    - VDC delete request.

Symptom: Intercloud Fabric and ICFC suffer a communication error while sending an ARP filter delete request to ICFC.

```
Disabling proxy ARP for Gateway (<enterprise gateway IP>) in ICS failed. Error: <message
 indicating type of ICFC communication exception>.
```

Possible Cause: Intercloud Fabric fails to send the ARP filter delete request to ICFC. This is due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:
   - Network delete request.
   - Network update request (dummy update).
   - VDC delete request.

Symptom: Intercloud Fabric and ICFC suffer a communication error while polling to review the outcome of an ARP filter delete operation in ICFC.

```
Disabling proxy ARP for Gateway (<enterprise gateway IP>) in ICS failed. Error: <message
 indicating type of ICFC communication exception>.
```

Possible Cause: Intercloud Fabric fails to receive the status of the Intercloud Fabric Switch instance to verify if the delete operation is successful or not. This is due to the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:
   - Network delete request.
   - Network update request (dummy update).
   - VDC delete request.

Symptom: Intercloud Fabric times out waiting for ICFC to delete the ARP filter.

```
Disabling proxy ARP for in ICS Gateway (<enterprise gateway IP>) in ICS failed. Error:
 <timeout message generated within Intercloud Fabric>.
```

Possible Cause: Intercloud Fabric gives up waiting for ICFC to create the ARP filter. ICFC suffers an internal error. This internal error prevents ICFC itself from timing out the ARP filter deletion flow.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to ICFC failing to complete the operation in time by restarting ICFC services.
3. Once the root cause is fixed, repeat the operation that initially triggered the interface deletion. The following are the possible operations:

    - Network delete request.
    - Network update request (dummy update).
    - VDC delete request.

Issues with the Intercloud Fabric Advanced Routing Service

This section includes symptoms, possible causes, and solutions for issues associated with the Advanced Routing Service.

Issues Enabling and Configuring Intercloud Fabric Advanced Routing Service

Symptom: The management network does not have two (both management and service) unallocated IP addresses.

```
Advanced Routing Service management and service IP address allocation from network
<network name> failed.
```

Possible Cause: Advanced Routing Service failed due to a lack of IP addresses in management network IP pool.

Verification and Solution:

1. Allocate at least two more IP addresses to the management network.
2. Delete the failed Intercloud Fabric link.
3. Create a new Intercloud Fabric link.

Symptom: The AWS marketplace End User License Agreement (EULA) not accepted for Advanced Routing Service images.

Possible Cause: The Advanced Routing Service instance creation fails due to EULA not being accepted for Advanced Routing Service images in AWS marketplace.

Verification and Solution:

1. Accept AWS EULA for Advanced Routing Service images.
2. Delete the failed Intercloud Fabric link.
3. Create a new Intercloud Fabric link.

Symptom: Intercloud Fabric and ICFC communication error.

```
Advanced Routing Service creation request to ICFC failed. <ICFC exception message>.
```

Possible Cause: The Advanced Routing Service creation fails due to an internal error within Intercloud Fabric. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, delete the Intercloud Fabric link.
4. Once successfully deleted, recreate the Intercloud Fabric link for the Advanced Routing Service.

Symptom: Intercloud Fabric times out during Advanced Routing Service creation.

```
Advanced Routing Service creation has timed out. ICFC internal error.
```

Possible Cause: Intercloud Fabric times out waiting for ICFC to create the Advanced Routing Service instance. This only occurs when ICFC suffers an internal error. This internal error prevents ICFC to timeout the Advanced Routing Service instance create flow.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Contact Cisco Customer Support to debug and fix the ICFC internal error.
3. Once successfully deleted, recreate the Intercloud Fabric link for the Advanced Routing Service.

Issues Disabling Intercloud Fabric Advanced Routing Service

Symptom: Intercloud Fabric and ICFC communication error.

Possible Cause: The Advanced Routing Service deletion fails due to an internal error within Intercloud Fabric. This can be due to one of the following reasons:

- Authentication failure.
- ICFC service is down.

Verification and Solution:

1. Review the Intercloud Fabric service request logs.
2. Fix the issue due to Intercloud Fabric failing to communicate with ICFC.
3. Once the root cause is fixed, if the Intercloud Fabric link deletion request fails, delete the Intercloud Fabric link again.

Symptom: Intercloud Fabric times out during the Advanced Routing Service delete attempt.

```
Undeploy Advanced Routing Service from Icflink <iclink name> has timed out: <timeout
message in ICFD>.
```

Possible Cause:

- Intercloud Fabric times out waiting for ICFC to delete the Advanced Routing Service instance.
- ICFC suffers an internal error.

Verification and Solution:

1. Contact Cisco Customer Support to debug and resolve the ICFC internal error.
2. Once the root cause is resolved, if the Intercloud Fabric link deletion request fails, delete the Intercloud Fabric link again.

Symptom: The Advanced Routing Service interface IP address release fails.

```
Failed to release IP <ip address> assigned to Advanced Routing Service instance <name>
 nic with role <management/service>.
```

```
Failed to release ips (<ip address(s)>) associated with Advanced Routing Service
interface <name>.
```

Possible Cause: Even if the deallocation of the resources within Intercloud Fabric fails, the Advanced Routing Service is successfully deleted. The allocation failure can occur for one or more of the following interface types:

- Data only Network interface
- Transport Network interface
- Management Network interface

Verification and Solution: Contact Cisco Customer Support to get the IP addresses released.

Issues Connecting Intercloud Fabric Advanced Routing Service to Networks

Symptom: The Intercloud Fabric link operational status is DOWN at the same time as the Advanced Routing Service create request is processing.

```
Advanced Routing Service <name> failed to create since link <name> is not-operational.
```

Possible Cause: Advanced Routing Service creation fails due to a network connection problem preventing Intercloud Fabric from reaching the Advanced Routing Service instance in the cloud for initial configuration:

- The operational status of the Intercloud Fabric link for the Intercloud Fabric cloud on which Advanced Routing Service is to be created must be UP before an attempt to create the Advanced Routing Service can be made.
- Once an Intercloud Fabric link is successfully deployed, Intercloud Fabric waits for a maximum of 15 minutes for the Intercloud Fabric link operational status to be UP.

Verification and Solution:

1. Resolve the Intercloud Fabric link operational DOWN status.
2. Delete the failed Intercloud Fabric link.
3. Create a new Intercloud Fabric link.

Note For more information on troubleshooting the Intercloud Fabric link, see Issues with the Intercloud Fabric Link.

Issues Configuring Intercloud Fabric Advanced Routing Service Static Route

Symptom: Failed to complete prerequisites for configuring Advanced Routing Service static route.

Possible Cause:

- If the Routing Service interface creation fails and is still pending, static route will not be added on Advanced Routing Service.
- Intercloud Fabric link status must be operational.
- Routing Service must be present.
- Advanced Routing Service must be successfully created.
- If static route configuration fails, static route status will be marked as FAILED_TO_CONFIGURE and corresponding service requests will go to a failed state.

Verification and Solution:

1. Ensure the Intercloud Fabric link status is operational.
2. Resubmit a network update on the corresponding network to retrigger the static route creation.

Symptom: Advanced Routing Service rejects the REST request to create, update, delete, and save the static routes.

Possible Cause: Review the SR logs to view the HTTP error from Advanced Routing Service.

Verification and Solution: Review the Advanced Routing Service troubleshooting steps for creation failure and then attempt to reconfigure the static route.