

High Availability

- Information About High Availability, on page 1
- Restrictions for High Availability, on page 5
- Configuring High Availability (GUI), on page 8
- Enabling High Availability (CLI), on page 10
- vWLC and N+1 High Availability, on page 12
- Adding a Hash Key to a Cisco vWLC (GUI), on page 13
- Adding a Hash Key to Cisco vWLC (CLI), on page 14
- Monitoring High Availability Standby Controller, on page 14
- Replacing the Primary Controller in an HA Setup, on page 16

Information About High Availability

High Availability in controllers allows you to reduce the downtime of the wireless networks that could occur due to failover of controllers.

A 1:1 (Active:Standby-Hot) stateful switchover of access points and clients is supported (HA SSO). In a High Availability architecture, one controller is configured as the primary controller and another controller as the secondary controller.

After you enable High Availability, the primary and secondary controllers are rebooted. During the boot process, the role of the primary controller is negotiated as active and the role of the secondary controller as standby-hot. After a switchover, the secondary controller becomes the active controller and the primary controller becomes the standby-hot controller. After subsequent switchovers, the roles are interchanged between the primary and the secondary controllers. The reason or cause for most switchover events is due to a manual trigger, a controller and/or a network failure.

During an HA SSO failover event, all of the AP CAPWAP sessions and client sessions in RUN state on the controller are statefully switched over to the standby controller without interruption, except PMIPv6 clients, which will need to reconnect and authenticate to the controller following an HA SSO switchover. For additional client SSO behaviors and limitations, see the "Client SSO" section in the *High Availability (SSO) Deployment Guide* at:

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-1/HA_SSO_DG/High_Availability_DG.html#pgfId-53637

The standby-hot controller continuously monitors the health of the active controller through its dedicated redundancy port. Both the controllers share the same configurations, including the IP address of the management interface.

Before you enable High Availability, ensure that both the controllers can successfully communicate with one another through their dedicated redundancy port, either through a direct cable connection or through Layer 2. For more details, see the "Redundancy Port Connectivity" section in the *High Availability (SSO) Deployment Guide*:

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-1/HA_SSO_DG/High_Availability_DG.html#pgfId-83028

In the Release 8.0 and later releases, the output of the **show ap join stats summary** command displays the status of the access points based on whether the access point joined the controller or it was synchronized from Active controller. One of the following statuses is displayed:

- Synched—The access point joined the controller before the SSO.
- Connected—The access point joined the controller after the SSO.
- Joined—The access point rejoined the controller, or a new AP has joined the controller after the SSO.

In Release 8.0 and later, the output of the **show redundancy summary** command displays the bulk synchronization status of access points and clients after the pair-up of active and standby controllers occurs. The values are:

- Pending— Indicates that synchronization of access points and the corresponding clients details from the active to standby controller is yet to begin.
- In-progress— Indicates that synchronization of access points and the corresponding clients details from the active to standby controller has begun and synchronization is in progress.
- Complete—Indicates that synchronization is complete and the standby controller is ready for a switchover to resume the services of the active controller.

From release 8.0 and later, in a High Availability scenario, the sleeping timer is synchronized between active and standby.

ACL and NAT IP configurations are synchronized to the High Availability standby controller when these parameters are configured before High Availability pair-up. If the NAT IP is set on the management interface, the access point sets the AP manager IP address as the NAT IP address.

The following are some guidelines for high availability:

- We recommend that you do not pair two controllers of different hardware models. If they are paired, the higher controller model becomes the active controller and the other controller goes into maintenance mode.
- We recommend that you do not pair two controllers on different controller software releases. If they are paired, the controller with the lower redundancy management address becomes the active controller and the other controller goes into maintenance mode.
- We recommend that you disable High Availability and add license in Cisco 5520and 8540 controllers (RTU based). However, it is not mandatory to disable High Availability as AP licenses added in Primary controller will be inherited to Secondary controller.
- All download file types, such as image, configuration, web-authentication bundle, and signature files are downloaded on the active controller first and then pushed to the standby-hot controller.
- Certificates should be downloaded separately on each controller before they are paired.

- You can upload file types such as configuration files, event logs, crash files, and so on, from the standby-hot controller using the GUI or CLI of the active controller. You can also specify a suffix to the filename to identify the uploaded file.
- To perform a peer upload, use the service port. In a management network, you can also use the redundancy management interface (RMI) that is mapped to the redundancy port or RMI VLAN, or both, where the RMI is the same as the management VLAN. Note that the RMI and the redundancy port should be in two separate Layer2 VLANs, which is a mandatory configuration.
- If the controllers cannot reach each other through the redundant port and the RMI, the primary controller becomes active and the standby-hot controller goes into the maintenance mode.



Note When the RMIs for two controllers that are a pair, and that are mapped to same VLAN and connected to same Layer3 switch stop working, the standby controller is restarted.

The `mobilityHaMac is out of range` XML message is seen during the active/standby second switchover in a High Availability setup. This occurs if mobility HA MAC field is more than 128.

- When High Availability is enabled, the standby controller always uses the Remote Method Invocation (RMI), and all the other interfaces—dynamic and management, are invalid.



Note The RMI is meant to be used only for active and standby communications and not for any other purpose.

- You must ensure that the maximum transmission unit (MTU) on RMI port is 1500 bytes or higher before you enable high availability.
- When High Availability is enabled, ensure that you do not use the backup image. If this image is used, the High Availability feature might not work as expected:
  - The service port and route information that is configured is lost after you enable SSO. You must configure the service port and route information again after you enable SSO. You can configure the service port and route information for the standby-hot controller using the **peer-service-port** and **peer-route** commands.
  - We recommend that you do not use the **reset** command on the standby-hot controller directly. If you use this, unsaved configurations will be lost.
- We recommend that you enable link aggregation configuration on the controllers before you enable the port channel in the infrastructure switches.
- All the configurations that require reboot of the active controller results in the reboot of the standby-hot controller.
- The Rogue AP Ignore list is not synchronized from the active controller to the standby-hot controller. The list is relearned through SNMP messages from Cisco Prime Infrastructure after the standby-hot controller becomes active.
- Client SSO related guidelines:

  - The standby controller maintains two client lists: one is a list of clients in the Run state and the other is a list of transient clients in all the other states.
  - Only the clients that are in the Run state are maintained during failover. Clients that are in transition, such as roaming, 802.1X key regeneration, web authentication logout, and so on, are dissociated.
  - As with AP SSO, Client SSO is supported only on WLANs. The controllers must be in the same subnet. Layer3 connection is not supported.

- In Release 7.3.x, AP SSO is supported, but client SSO is not supported, which means that after a High Availability setup that uses Release 7.3.x encounters a switchover, all the clients associated with the controller are deauthenticated and forced to reassociate.
- You must manually configure the mobility MAC address on the then active controller post switchover, when a peer controller has a controller software release that is prior to Release 7.2.
- To enable an access point to maintain controlled quality of service (QoS) for voice and video parameters, all the bandwidth-based or static call admission control (CAC) parameters are synchronized from active to standby when a switchover occurs.
- The standby controller does not reboot; instead enters the maintenance mode when unable to connect to the default gateway using the redundant port. Once the controller reconnects to the default gateway, the standby controller reboots and the High Availability pair with the active controller is initiated. However, the active controller still reboots before entering the maintenance mode.
- The following are supported from Release 8.0:
  - Static CAC synchronization—To maintain controlled Quality-of-Service (QoS) for voice and video parameters, all the bandwidth-based or static CAC parameters services are readily available for clients when a switchover occurs.
  - Internal DHCP server—To serve wireless clients of the controller, the internal DHCP server data is synchronized from the active controller to the standby controller. All the assigned IP addresses remain valid, and IP address assignation continues when the role changes from active to standby occurs.
  - Enhanced debugging and serviceability—All the debugging and serviceability services are enhanced for users.

- The physical connectivity or topology of the access points on the switch are not synchronized from the active to the standby controller. The standby controller learns the details only when the synchronization is complete. Hence, you must execute the **show ap cdp neighbors all** command only after synchronization is complete, and only when the standby becomes the then active controller.
- To enable access points to join the HA SKU secondary controller that has been reset to factory defaults, you must:
  - Configure the HA SKU controller as secondary controller. To do this, you must run the **config redundancy unit secondary** command on the HA SKU controller.
  - Reboot the HA SKU controller after you successfully execute the **config redundancy unit secondary** command.

Redundancy Management Interface

The active and standby-hot controllers use the RMI to check the health of the peer controller and the default gateway of the management interface through network infrastructure.

The RMI is also used to send notifications from the active controller to the standby-hot controller if a failure or manual reset occurs. The standby-hot controller uses the RMI to communicate to the syslog, NTP/SNTP server, FTP, and TFTP server.

It is mandatory to configure the IP addresses of the Redundancy Management Interface and the Management Interface in the same subnet on both the primary and secondary controllers.

Redundancy Port

The redundancy port is used for configuration, operational data synchronization, and role negotiation between the primary and secondary controllers.

The redundancy port checks for peer reachability by sending UDP keepalive messages every 100 milliseconds (default frequency) from the standby-hot controller to the active controller. If a failure of the active controller occurs, the redundancy port is used to notify the standby-hot controller.

If an NTP/SNTP server is not configured, the redundancy port performs a time synchronization from the active controller to the standby-hot controller.

The redundancy ports can connect over an L2 switch. Ensure that the redundancy port round-trip time is less than 80 milliseconds if the keepalive timer is set to default, that is, 100 milliseconds, or 80 percent of the keepalive timer if you have configured the keepalive timer in the range of 100 milliseconds to 400 milliseconds. The failure detection time is calculated, for example, if the keepalive timer is set to 100 milliseconds, as follows: 3 * 100 = 300 + 60 = 360 + jitter (12 milliseconds) = ~400 milliseconds. Also, ensure that the bandwidth between redundancy ports is 60 Mbps or higher. Ensure that the maximum transmission unit (MTU) is 1500 bytes or higher.

Related Documentation

- *High Availability (SSO) Deployment Guide*—https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-1/HA_SSO_DG/High_Availability_DG.html
- *N+1 High Availability Deployment Guide*—https://www.cisco.com/c/en/us/td/docs/wireless/technology/hi_avail/N1_High_Availability_Deployment_Guide.html

Restrictions for High Availability

- We recommend that you do not disable LAG physical ports when HA SSO is enabled.
- HA sync for Fabric-related statistics is not supported.
- You should apply an access list for SSH to the redundancy interface on upper switch, if controller is configured for HA SSO and redundancy management is configured over a dynamic interface. Failure to do so enables the SSH client to connect through the redundancy management interface regardless of the CPU ACL.
- In an HA environment using FlexConnect locally switched clients, the client information might not show the username. To get details about the client, you must use the MAC address of the client. This restriction does not apply to FlexConnect centrally switched clients or central (local) mode clients.

- In an HA environment, an upgrade from an LDPE image to a non-LDPE image is not supported.
- It is not possible to pair two primary controllers or two secondary controllers.
- Standby controllers are unavailable on the APs connected switch port.
- An HA-SKU controller with an evaluation license cannot become a standby controller. However, an HA-SKU controller with zero license can become a standby controller.
- In an HA setup, CPU-ACL cannot be applied on the service port. However, if you want to block the service port using CPU-ACL, you can use the command **config acl high-priority** to configure as required.
- Service VLAN configuration is lost when moving from HA mode to non-HA mode and conversely. You should then configure the service IP address manually again.
- The following scenario is not supported: The primary controller has the management address and the redundancy management address in the same VLAN, and the secondary controller has the management address in the same VLAN as the primary one, and the redundancy management address in a different VLAN.
- The following is a list of some software upgrade scenarios:
  - A software upgrade on the active controller ensures the upgrade of the standby-hot controller.
  - An in-service upgrade is not supported. Therefore, you should plan your network downtime before you upgrade the controllers in an HA environment.
  - Rebooting the active controller after a software upgrade also reboots the standby-hot controller.
  - We recommend that both active and standby-hot controllers have the same software image in the backup before running the **config boot backup** command. If both active and standby-hot controllers have different software images in the backup, and if you run the **config boot backup** command in the active controller, both the controllers reboot with their respective backup images breaking the HA pair due to a software mismatch.
  - A schedule reset applies to both the controllers in an HA environment. The peer controller reboots a minute before the scheduled time expires on the active controller.
  - You can reboot the standby-hot controller from the active controller by entering the **reset peer-system** command if the scheduled reset is not planned. If you reset only the standby-hot controller with this command, any unsaved configurations on the standby-hot controller are lost. Therefore, ensure that you save the configurations on the active controller before you reset the standby-hot controller.
  - If an SSO is triggered at the time of the image transfer, a preimage download is reinitiated.
  - Only **debug** and **show** commands are allowed on the standby-hot controller.
  - After a switchover, if a peer controller has a controller software release that is prior to Release 7.5, all the mobility clients are deauthenticated.
- It is not possible to access the standby-hot controller through the controller GUI, Cisco Prime Infrastructure, or Telnet. You can access the standby-hot controller only on its console.
- When you enable both RADIUS profiling and WSA in an SSID, local profiling gets enabled in the same SSID.
- In an HA setup, client Tx or Rx packets are not sent to the standby controller, hence, Remote Method Invocation (RMI) is not supported.

- When a failover occurs, the standby controller must be in a standby-hot state and the redundant port in a terminal state in SSO for successful switchover to occur.
- To enable or disable LAG, you must disable HA.



Note If LAG is disabled and both primary and backup ports are connected to the management interface and if the primary port becomes nonoperational, a switchover might occur because the default gateway is not reachable and backup port failover might exceed 12 seconds.

- When a failover occurs and the standby controller becomes the new active controller, it takes approximately 15–20 minutes to synchronize the database (AP, client, and multicast) between the two controllers. If another failover occurs during this time, the HA structures would not yet be synchronized. Therefore, the APs and clients would have to get reassociated and reauthenticated respectively.
- Pairwise Master Key (PMK) cache synchronization is not supported on FlexConnect local-authenticated clients.
- Client SSO restrictions:
  - New mobility is not supported.
  - Posture and network admission control out-of-band are not supported because the client is not in the Run state.
  - The following are not synchronized between the active and standby controller:
    - Cisco Compatible Extension-based applications
    - Client statistics
    - Proxy Mobile IPv6, Application Visibility and Control, session initiation protocol (SIP), and static call admission control (CAC) tree
    - Workgroup bridges and the clients that are associated with them
    - Passive clients
  - Encryption is supported.
- Encryption is supported only if the active and standby controllers communicate through the Redundancy Management Interface on the management ports. Encryption is not supported if the redundancy port is used for communication between the active and standby controllers.
- You cannot change the NAT address configuration of the management interface when the controllers are in redundancy mode. To enable NAT address configuration on the management interface, you must remove the redundancy configuration first, make the required changes on the primary controller, and then reenable the redundancy configuration on the same controller.
- After you enable SSO, you must access both the standby and active controller using:
  - The console connection
  - SSH facility on the service port

- SSH facility on the redundant management interface

- Synchronization of bulk configurations is supported only for the configurations that are stored in XMLs. Scheduled reboot is a configuration that is not stored in XMLs or Flash. Therefore, the scheduled reboot configuration is not included in the synchronization of bulk configurations.
- When a switchover occurs, the controller does not synchronize the information on DHCP dirty bit from the active to standby controller even when DHCP dirty bit is set on the active controller. After a switchover, the controller populates the DHCP dirty bit based on the client DHCP retries.

Configuring High Availability (GUI)

Before you begin

Ensure that the management interfaces of both controllers are in the same subnet. You can verify this on the GUI of both the controllers by choosing **Controllers** > **Interfaces** and viewing the IP addresses of the management interface.

Procedure

Step 1 On the GUI of both the controllers, choose **Controller** > **Redundancy** > **Global Configuration**.

The **Global Configuration** window is displayed.

Step 2 Enter the addresses of the controllers in the **Redundant Management IP** field and the **Peer Redundant Management IP** field.

Note
Ensure that the Redundant Management Interface IP address of one controller is the same as the Redundant Management Interface IP address of the peer controller.

Step 3 From the **Redundant Unit** drop-down list, choose one of the controllers as primary and the other as secondary.

Step 4 On the GUI of both the controllers, set the **SSO** to **Enabled** state.

Note
After you enable an SSO, the service port peer IP address and the service port netmask appear on the configuration window. Note that the service port peer IP address and the netmask can be pushed to the peer only if the HA peer is available and operational. When you enable HA, you do not have to configure the service port peer IP address and the service port netmask parameters. You must configure the parameters only when the HA peer is available and operational. After you enable SSO, both the controllers are rebooted. During the reboot process, the controllers negotiate the redundancy role through the redundant port, based on the configuration. The primary controller becomes the active controller and the secondary controller becomes the standby controller.

Step 5 [Optional] After the HA pair becomes available and operational, you can configure the peer service port IP address and the netmask after the service port is configured as static. If you enable DHCP on the service port, you do not have to configure these parameters on the **Global Configuration** window:

- **Service Port Peer IP**—IP address of the service port of the peer controller.

- **Service Port Peer Netmask**—Netmask of the service port of the peer controller.
- **Mobility MAC Address**—A common MAC address for both the active and standby controllers that is used in the mobility protocol. If an HA pair has to be added as a mobility member for a mobility group, the mobility MAC address (instead of the system MAC address of the active or standby controller) should be used. Normally, the mobility MAC address is chosen as the MAC address of the active controller and you do not have to manually configure this.
- **Keep Alive Timer**—The timer that controls how often the standby controller sends keepalive messages to the active controller. The valid range is between 100 to 1000 milliseconds.
- **Peer Search Timer**—The timer that controls how often the active controller sends peer search messages to the standby controller. The valid range is between 60 to 300 seconds.

Note
After you enable the HA and pair the controllers, there is only one unified GUI to manage the HA pair through the management port. GUI access through the service port is not feasible for both the active and standby controllers. The standby controller can be managed only through the console port or the service port.

Only Telnet and SSH sessions are allowed through the service port of the active and standby controllers.

Step 6 Click **Save Configuration**.

Step 7 View the redundancy status of the HA pair by choosing **Monitor > Redundancy > Summary**.

The **Redundancy Summary** window is displayed.

Step 8 View the redundancy status of the HA pair by choosing **Monitor > Redundancy > Detail**.

The **Redundancy Detail** page is displayed.

Step 9 View the redundancy statistics information of the HA pair by choosing **Monitor > Redundancy > Statistics**.

The **Redundancy Statistics** page is displayed.

Step 10 (Optional) Perform these steps to configure the peer network route:

a) Choose **Controller** > **Redundancy** > **Peer Network Route**.

The **Network Routes Peer** window is displayed.

This window provides a summary of the existing service port network routes of the peer controller to network or element management systems on a different subnet. You can view the IP address, IP netmask, and gateway IP address.

b) To create a new peer network route, click **New**.
c) Enter the **IP address**, **IP netmask**, and the **Gateway IP address** of the route.
d) Click **Apply**.

Enabling High Availability (CLI)

Procedure

Step 1 Before you configure HA, it is mandatory to have the management interface of both the controllers in the same subnet. See the interface summary information by entering these commands on both the controllers:

show interface summary

Step 2 High Availability is disabled by default. Before you enable HA, it is mandatory to configure the redundancy management IP address and the peer redundancy management IP address. Both the interfaces must be in the same subnet as the management interface. Enter the following commands to configure the redundancy management IP addresses:

- On WLC1: **config interface redundancy-management** *redundancy-mgmt-ip-addr-wlc1* **peer-redundancy-management** *peer-redundancy-mgmt-ip-addr-wlc2*
- On WLC2: **config interface redundancy-management** *redundancy-mgmt-ip-addr-wlc2* **peer-redundancy-management** *peer-redundancy-mgmt-ip-addr-wlc1*

Step 3 Configure one controller as primary (by default, the controller HA Unit ID is primary and should have a valid AP-BASE count license installed) and another controller as secondary (AP-BASE count from the primary controller is inherited by this unit) by entering these commands:

- WLC1 as primary—**config redundancy unit primary**
- WLC2 as secondary—**config redundancy unit secondary**

Note
You are not required to configure the unit as secondary if it is a factory-ordered HA SKU that can be ordered from Release 7.3 onwards. A factory-ordered HA SKU is a default secondary unit and takes the role of the standby controller the first time it is paired with an active controller that has a valid AP count license.

Step 4 After you have configured the controllers with redundancy management and peer redundancy management IP addresses and have configured the redundant units, you must enable SSO. Ensure that the physical connections are operational between both the controllers (that is, both the controllers are connected back to back via the redundant port using an Ethernet cable) and the uplink is also connected to the infrastructure switch and the gateway is reachable from both the controllers before SSO is enabled.

After SSO is enabled, controllers are rebooted. During the boot process, the controllers negotiate the HA role as per the configuration via the redundant port. If the controllers cannot reach each other via the redundant port or via the redundant management interface, the controller that is configured as secondary might go into maintenance mode.

Enable SSO on both the controllers by entering these commands:

config redundancy mode sso

Note
Enabling SSO initiates a controller reboot.

Step 5 Enabling SSO reboots the controllers to negotiate the HA role as per the configuration performed. Once the role is determined, configuration is synchronized from the active controller to the standby controller via the redundant port. Initially, the controller configured as secondary reports XML mismatch and downloads the configuration from the active controller and reboot again. During the next reboot after determining the HA role, the controller validates the configuration again, reports no XML mismatch, and process further to establish itself as the standby controller.

Note
Once SSO is enabled, you can access the standby controller through a console connection or through SSH on the service port and on the redundant management interface.

Step 6 After SSO is enabled, controllers are rebooted, the XML configuration is synchronized, WLC1 transitions its state to active and WLC2 transitions its state to standby hot. From this point, GUI, Telnet, and SSH for WLC2 on the management interface does not work because all the configurations and management must be done from the active controller. If required, the standby controller (WLC2) can be managed only through the console or service port.

Once the peer controller transitions to the standby hot state, the *Standby* keyword is automatically appended to the standby controller's prompt name.

Step 7 To see the redundancy summary information for both the controllers, enter this command:

show redundancy summary

Configuring High Availability Parameters (CLI)

Procedure

- Configure the IP address and netmask of the peer service port of the standby controller by entering this command:

  config redundancy interface address peer-service-port *ip-address netmask*

  This command can be run only if the HA peer controller is available and operational.

- (Optional) Configure the route configurations of the standby controller by entering this command:

  config redundancy peer-route {**add** *network-ip-addr ip-mask* | **delete** *network-ip-addr*}

Note This command can be run only if the HA peer controller is available and operational.

- (Optional) Configure a mobility MAC address by entering this command:

  config redundancy mobilitymac *mac-addr*

Note

- This command can be run only when SSO is disabled.
- From Release 8.0.132.0 onwards, mobility MAC configuration is no longer present in the uploaded configuration. Therefore, if you download this configuration file back to the controller, you must add the **config redundancy mobilitymac** *mac_address* command in the config file before download.

- Configure a redundancy timer by entering this command:

  config redundancy timer {**keep-alive-timer** *time-in-milliseconds* | **peer-search-timer** *time-in-seconds*}

- View the status of the redundancy by entering this command:

  show redundancy {**summary** | **detail**}

- View information about the redundancy management interface by entering this command:

  show interface detailed redundancy-management

- View information about the redundancy port by entering this command:

  show interface detailed redundancy-port

- Reboot a peer controller by entering this command:

  reset peer-system

- Start the upload of file types, such as configuration, event logs, crash files, and so on from the standby-hot controller by entering this command on the active controller:

  transfer upload peer-start

- View information about sleeping clients after a switchover, by entering this command on the then active controller :

  show custom-web sleep-client summary

vWLC and N+1 High Availability

Release 8.4 introduces support for N+1 High Availability (HA) on the Cisco Virtual Wireless Controller (vWLC) platform. For information on how to configure HA, see:

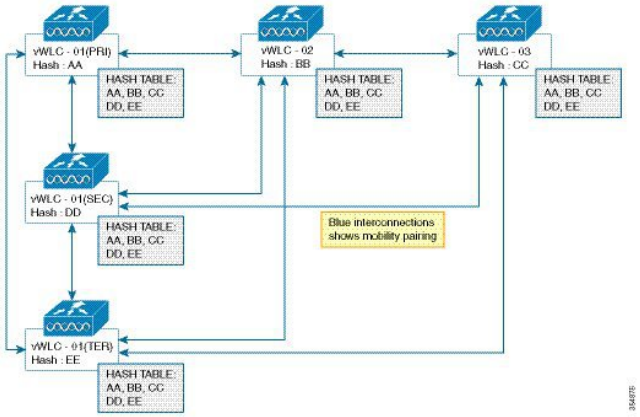
https://www.cisco.com/c/en/us/td/docs/wireless/technology/hi_avail/N1_High_Availability_Deployment_Guide/N1_HA_Overview.html#pgfId-1054644

The Cisco vWLC HA has the following prerequisites:

- The primary, secondary, and tertiary vWLCs should be part of the same mobility group.
- The vWLC in the mobility group should have a uniform set of hash keys to seamlessly move an AP from one vWLC to another. For example, if we have vWLCs, N, in a mobility group, or vWLC, M, and normal controllers (where M is greater than N), then all vWLCs should have the hashes of other vWLCs in the same group.
- For effective connectivity of the APs on all the vWLCs in a mobility group (including vWLC mobility members in N+1 format), the mobility hash table should contain all the vWLC hash keys.

Note A hash table works only when vWLCs are paired as mobility members.

Figure 1: vWLC N+1 in a Mobility Group



Adding a Hash Key to a Cisco vWLC (GUI)

Perform the procedure given below to add hash key to Cisco vWLC.

Before you begin

Create mobility peers before adding a hash key to Cisco vWLC.

Procedure

Step 1 Choose **Controller** > **Mobility Management** > **Mobility Groups**.

The **Static Mobility Group Members** window displays the existing members and the hash keys configured for them.

Step 2 Click **New**.

The **Mobility Group Member** > **New** window is displayed.

Step 3 In the **Member IP Address(Ipv4/Ipv6)** field, enter the member's IP address. In the **Member MAC Address** field, enter the member's MAC address. In the **Group Name** field, enter the group name. In the **Hash** field, enter the hash key.

Step 4 Click **Apply**.

Adding a Hash Key to Cisco vWLC (CLI)

Perform the procedure given below to add a hash key to a Cisco vWLC, using the CLIs.

- Read the hash key.
- Copy the hash key to the other members of the mobility group.
- Verify the mobility hash configuration.

Before you begin

- The hash value should be unique for each vWLC.
- Create mobility peers before adding a hash key to a vWLC.

Procedure

Step 1 **show mobility group member hash**

Example:

```
(Cisco Controller)> show mobility group member hash
```

Reads the existing hash key.

Step 2 **config mobility group member hash** *ipv4-address hash-key*

Example:

```
(Cisco Controller)> config mobility group member hash 9.11.34.55
1f81d80082e9d30312d3b4920be22aed34b93b56
```

Copies the hash to other members in the mobility group.

Step 3 **show mobility group member hash**

Example:

```
(Cisco Controller)> show mobility group member hash
Default Mobility Domain..................... default

 IP Address      Hash Key
------------------------------------------------------------
 9.11.34.55     1f81d80082e9d30312d3b4920be22aed34b93b56
```

Verifies the mobility hash configuration on all the mobility members in the group.

Monitoring High Availability Standby Controller

You can view the status and health information of active and standby controller separately. This section describes the details of getting health information and traps from the standby controller.

The standby controller uses the redundancy management interface for any external communications such as when talking to Syslog, NTP server, TFTP server, and so on. On the standby controller, the management user authentication and accounting is performed on the redundancy management interface. RADIUS or TACACS+ server can be used for user authentication, apart from a local management user account. To support this, the redundancy interface IP address(es) should be added as network device on the RADIUS or TACACS+ server. The authentication request is sent to RADIUS or TACACS+ server over redundancy management interface. Whenever you log on to the standby controller, accounting message is sent to the RADIUS server. The purpose of the accounting message is to log the admin logon events on the standby controller console.

This feature is supported on all controller models supporting HA SSO feature:

- Cisco 8500 Series Controllers
- Cisco 3504 Controllers
- Cisco 5500 Series Controllers

Events and Notifications

- Trap when controller becomes Hot Standby—A trap is reported with time stamp when HA peer becomes Hot Standby and the trap shown below is reported

  "RF notification EventType:37 Reason :HA peer is Hot-Standby...At:..."

  A new trap type is added in CISCO-RF-SUPPLEMENTAL-MIB.my

- Trap when Bulk Sync Complete—After the HA pairing is done and Bulk sync is complete, the following trap is reported:

  "RF notification EventType:36 Reason :Bulk Sync Completed...At:.."

  A new trap type is added in CISCO-RF-SUPPLEMENTAL-MIB.my

- Trap when Standby controller goes down—When the standby peer goes down due to manual reset, crash, memory leak/hang, or moving to maintenance mode, the following trap is reported:

  "RF failure notification ErrorType: 34 Reason :Lost Peer, Moving to Active-No-Peer State!"

  On the CLI, you can view the trap by entering the **show traplog** command.

- Syslog notification when Admin login on Standby

  1. Admin login to Standby via SSH generates an event in msglog/syslog. The following is a sample system message:

     ***emWeb: Mar 06 20:34:42.675: #CLI-3-LOGIN_STANDBY: [SS] cli_lvl7.c:4520 [USER@9 name="admin" from="SSH"] user login success on standby controller.**

     You can view this message on the standby controller by entering the **show msglog** command.

  2. Admin login to Standby via console generates an event in msglog/syslog. The following is a sample system message:

     ***emWeb: Mar 06 20:34:42.675: #CLI-3-LOGIN_STANDBY: [SS] cli_lvl7.c:4520 [USER@9 name="admin" from="console"] user login success on standby controller.**

     You can view this message on the standby controller by entering the **show msglog** command.

- Peer Process Statistics—The CPU and Memory statistics of all the threads of the standby controller are synchronized with the active controller every 10 seconds. This information is displayed when you query for the Peer statistics on the active controller.

  Enter these commands on the active controller to view the peer process system, CPU, and memory statistics:

  - **show redundancy peer-system statistics**
  - **show redundancy peer-process cpu**
  - **show redundancy peer-process memory**

  On the GUI, choose **Monitor** > **Redundancy** > **Peer Statistics** to view the peer process system, CPU, and memory statistics.

Replacing the Primary Controller in an HA Setup

In an HA setup, suppose the primary controller is not operational and you are required to replace it; the standby controller is operational with all the APs associated with it; and the new controller received return material authorization (RMA) that can be added with one of the failed controllers in the HA pair. Follow these steps to replace the primary controller in an active HA setup:

Procedure

Step 1 Ensure that the new controller and the controller to be replaced are running the same version of the controller software.

Step 2 Configure the new controller with the same subnet management IP addresses as the controller to replaced.

Step 3 Configure the new controller with HA configuration that includes redundancy management, IP address, and peer primary. Accept the licensing EULA on the primary controller and then enable AP SSO.

Note
If you enable AP SSO without accepting the EULA, the controllers do not synchronize.

Step 4 When AP SSO is enabled, the controller reboots. While the controller reboots, the AP SSO discovers the currently active standby controller, synchronized the configuration, and transitions to a standby-hot state.

Note
You do not need to break the HA configuration on the current active controller or reboot the current active controller. The configuration will be synchronized with the current active controller.