



# Troubleshooting Mesh Access Points

- [Installation and Connections](#), on page 1

## Installation and Connections

- Step 1** Connect the mesh access point that you want to be the RAP to the controller.
- Step 2** Deploy the radios (MAP) at the desired locations.
- Step 3** On the controller CLI, enter the **show mesh ap summary** command to see all MAPs and RAPs on the controller.

*Figure 1: Show Mesh AP Summary Page*

```
(Cisco Controller) >show mesh ap summary
```

AP Name	AP Model	BVI MAC	CERT MAC	Hop	Bridge Group Name	Enhanced Feature Set
1532MAP2-DaisyChained	AIR-CAP1532E-A-K9	4c:4e:35:46:f2:72	4c:4e:35:46:f2:72	0	default	N/A
1532RAP1	AIR-CAP1532E-A-K9	4c:4e:35:46:f2:64	4c:4e:35:46:f2:64	0	default	N/A
1532MAP1	AIR-CAP1532E-A-K9	4c:4e:35:46:f1:4e	4c:4e:35:46:f1:4e	1	default	N/A
1524PSRAP1	AIR-LAP1524PS-A-K9	00:22:be:41:23:00	00:22:be:41:23:00	0	MESHDEM01	N/A
1522MAP2	AIR-LAP1522AG-A-K9	00:22:be:42:fe:00	00:22:be:42:fe:00	1	MESHDEM01	N/A

  

```
Number of Mesh APs..... 3  
Number of RAPs..... 2  
Number of MAPs..... 1  
Number of Flex+Bridge APs..... 2  
Number of Flex+Bridge RAPs..... 1  
Number of Flex+Bridge MAPs..... 1
```

- Step 4** On the controller GUI, click **Wireless** to see the mesh access point (RAP and MAP) summary.

Figure 2: All APs Summary Page

All APs

Search by AP MAC

AP Name	AP MAC	AP Up Time	Admin Status	Operational Status	AP Mode	Certificate Type
<a href="#">iMeshRap1</a>	00:19:30:76:32:72	0 d, 22 h 24 m 25 s	Enable	REG	Local	MIC
<a href="#">HJRAP1</a>	00:1d:71:0d:e1:00	0 d, 22 h 12 m 37 s	Enable	REG	Bridge	MIC
<a href="#">HJMAP3</a>	00:1d:71:0d:d5:00	0 d, 22 h 05 m 04 s	Enable	REG	Bridge	MIC
<a href="#">HJMAP1</a>	00:1d:71:0c:f4:00	0 d, 22 h 04 m 48 s	Enable	REG	Bridge	MIC
<a href="#">HJMAP2</a>	00:1d:71:0c:f0:00	0 d, 22 h 04 m 53 s	Enable	REG	Bridge	MIC
<a href="#">HPRAP1</a>	00:1e:14:48:43:00	0 d, 05 h 35 m 24 s	Enable	REG	Bridge	MIC
<a href="#">HPMAP1</a>	00:1b:d4:a7:78:00	0 d, 22 h 04 m 25 s	Enable	REG	Bridge	MIC

273952

**Step 5** Click **AP Name** to see the details page and then select the **Interfaces** tab to see the active radio interfaces.

The radio slot in use, radio type, subband in use, and operational status (UP or DOWN) are summarized.

- All APs supports 2 radio slots: slot 0—2.4 GHz and slot 1—5 GHz.

If you have more than one controller connected to the same mesh network, then you must specify the name of the primary controller using global configuration for every mesh access point or specify the primary controller on every node, otherwise the least loaded controller is the preferred controller. If the mesh access points were previously connected to a controller, they already have learned a controller’s name.

After configuring the controller name, the mesh access point reboots.

**Step 6** Click **Wireless > AP Name** to check the mesh access point’s primary controller on the AP details page.

## Debug Commands

The following two commands are very helpful to see the messages being exchanged between mesh access points and the controller.

```
(Cisco Controller) > debug capwap events enable
(Cisco Controller) > debug disable-all
```

You can use the **debug** command to see the flow of packet exchanges that occur between the mesh access point and the controller. The mesh access point initiates the discovery process. An exchange of credentials takes place during the join phase to authenticate that the mesh access point is allowed to join the mesh network.

Upon a successful join completion, the mesh access point sends a CAPWAP configuration request. The controller responds with a configuration response. When a Configure Response is received from the controller, the mesh access point evaluates each configuration element and then implements them.

## Remote Debug Commands

You can log on to the mesh access point console for debugging either through a direct connection to the AP console port or through the remote debug feature on the controller.

To invoke remote debug on the controller, enter the following commands:

```
(Cisco Controller) > debug ap enable ap-name
(Cisco Controller) > debug ap command command ap-name
```

## AP Console Access

AP1500s have a console port. A console cable is not shipped with the mesh access point. For the 1550 series access points, console ports are easily accessible and you need not open the access point box.

The AP1500s have console access security embedded in the code to prevent unauthorized access on the console port and provide enhanced security.

The **login ID** and **password** for console access are configured from the controller. You can use the following commands to push the username/password combination to the specified mesh access point or all access points:

```
<Cisco Controller> config ap username cisco password cisco ?
all          Configures the Username/Password for all connected APs.
<Cisco AP>  Enter the name of the Cisco AP.

<Cisco Controller> config ap username cisco password cisco all
```

You must verify whether the username/password pushed from the controller is used as *user-id* and *password* on the mesh access point. It is a nonvolatile setting. Once set, a *login ID* and *password* are saved in the private configuration of the mesh access point.

Once you have a successful login, the trap is sent to the Cisco Prime Infrastructure. If a user fails to log on three times consecutively, login failure traps are sent to the controller and Cisco Prime Infrastructure.



**Caution**

A mesh access point must be reset to the factory default settings before moving from one location to another.

**Hardware Reset**

Perform a hardware reset on this AP

Reset AP Now

**Set to Factory Defaults**

Clear configuration on this AP and reset it to factory defaults

Clear Config

206711

## Cable Modem Serial Port Access From an AP

Commands can be sent to the cable modem from the privileged mode of the CLI. Use the command to take a text string and send it to the cable modem UART interface. The cable modem interprets the text string as one of its own commands. The cable modem response is captured and displayed on the Cisco IOS console. Up to 9600 characters are displayed from the cable modem. Any text that is greater than 4800 characters is truncated.

The modem commands are only operational on mesh APs that have devices connected to the UART port originally intended for the cable modem. If the commands are used on a mesh AP that does not have a cable modem (or any other device connected to the UART), the commands are accepted, however, but they do not produce any returned output. No errors are explicitly flagged.

## Configuration

Enter the following command from the privileged mode of the MAP:

```
AP#send cmodem timeout-value modem-command
```

The modem command is any command or text to send to the cable modem. The range of timeout value is 1 to 300 seconds. However, if the captured data equals 9600 characters, any text beyond that is truncated and the response, irrespective of the timeout value and is immediately displayed on the AP console.

**Figure 3: Cable Modem Console Access Command**

```
RAP-CM-N1#send ?
*          All tty lines
<0-16>    Send a message to a specific line
cmodem    Enter cable modem command
console   Primary terminal line
log       Logging destinations
vty      Virtual terminal

RAP-CM-N1#send cmodem ?
LINE      Enter modem command string
<cr>
```

279059

Figure 4: Cable Modem Console Access Command

```

RAP-CM-N1#send cmodem ls
ls
CM>
CM> ls

!                ?                REM                cd                dir
find_command    help                history            instances        ls
man             pwd                sleep             syntax           system_time
usage
----
mbufShow        memShow            mutex_debug       ping             read_memory
reset           routeShow         run_app           shell            stackShow
start_idle_profiling  stop_idle_profiling  taskDelete
taskInfo        taskPrioritySet   taskResume       taskShow        taskSuspend
taskTrace       usfsShow          version           write_memory     zone
----
[HeapManager] [SA] [cm_hal] [docsis_ctl] [embedded_target] [enet_hal]
[event_log] [flash] [forwarder] [ip_hal] [msgLog] [non-vol] [pingHelper]
[snmp] [snoop] [usb_hal]

CM>
RAP-CM-N1#send cmodem cd docsis
cd
CM>
CM> cd docsis
CM> cd docsis

Active Command Table:  CM DOCSIS Control Thread Commands (docsis_ctl)

CM -> docsis_ctl

CM/DocsisCtl>
RAP-CM-N1#
    
```



**Caution** The question mark (?) and the exclamation point (!) should not be used in the **send cmodem** command. These characters have immediate interpreted use in the Cisco IOS CLI. Therefore, they cannot be sent to the modem.

**Enabling the Cable Modem Console Port**

By default, the Cable Modem console port is disabled. This is to prevent users from accessing the console through their residential cable modem. In the AP1572IC, AP1572EC, and AP1552C model, the cable modem console is connected directly to the access point. The console port is required for signaling between the AP and the cable modem. There are two methods to enable the cable modem console port, either through SNMP or by adding the command to the configuration .cm file on the CMTS.



**Note** For the AP1572EC, AP1572IC, AP1552C, and AP1552CU, the cable modem must be enabled.

- Enable the cable modem console port through SNMP by entering this command to the IP address of the cable modem:

```
snmpset -c private IP_ADDRESS cmConsoleMode.0 i N
```

Using the OID, enter this command:

```
snmpset -c private IP_ADDRESS
1.3.6.1.4.1.1429.77.1.4.7.0 i N
```

Where IP\_ADDRESS is any IPv4 address and N is an integer, 2 to enable read-write, 1 for read-only, or 0 to disable.

Example:

```
snmpset -c private 209.165.200.224 cmConsoleMode.0 i 2
```

- Enable the cable modem console port through the configuration file. The configuration file (with a .cm extension) is loaded into the cable modem head end. It is pushed to the cable modem as part of the join process. Enter the following line to the cable modem configuration file:

```
SA-CM-MIB::cmConsoleMode.0 = INTEGER: readWrite(2)
```

Using the OID, enter this line:

```
SA-CM-MIB::cmConsoleMode.0 = INTEGER: readWrite(2)
```

### Resetting the AP1572xC/AP1552C Through the Cable Modem

An AP can be reset by entering an SNMP command to the Cable Modem, which resides inside the access point. For this feature to work, you must enable the cable modem console port.

Reset the AP by entering this snmpset command:

```
Snmpset -v2c -c public IP ADDRESS 1.3.6.1.4.1.1429.77.1.3.17.0 i 1
```

Where the IP ADDRESS is the IPv4 address of the cable modem.

## Mesh Access Point CLI Commands

You can enter these commands directly on the mesh access point using the AP console port or you can use the remote debug feature from the controller:

```

H1 #show mesh ?
  adjacency      l'ESH Adjacency
  astools        l'ESH Anti-strand tools
  backhaul       l'ESH backhaul
  channel        l'ESH channel
  canfig         l'ESH config parameter
  dfs            l'ESH dfs information
  ethernet       show mesh Ethernet bridging
  forwarding     l'ESH Forwarding
  inventory      platform inventory
  linktest       l'ESH linktest stats
  module        l'ESH module detail
  nplrf         l'ESH NBN tool
  security       l'ESH Security show
  simulation     show simulated configuration
  status         l'ESH status
  
```

```

H1 #show mesh config
rtsfhreshold1 la 0, ehs 0, a.1lin 0, co.1lex 0
rtsfhreshold1 lb 0, aifs 0, a.1Hin 0, a.1lax 0
huRetries 0. 1lri <Rate 0 qDepth 0
802.11M Client Statistics Push Int. .... al: 3
range parameter: 12000
mesh security node: 0
Universal Client Access: disabled
public safety global state: enabled
Battery backup state: enabled
multicast node: in-out
Full Sector DFS: enabled
  
```

```
HJRAP111lehou capl01Bp client mb
AdminState          ADHIN ENABLED
SuVer               S. 2.98.0
NumFl1 ledSlots    2
Name              HJRAP1
Location         default location
Huarllame           SEYf-CliffROLLER
Huarrlp             209.165.200.227
Huartt.Ner         0.0.0.0
ApHocle             Brl d!JE!
ApSubl'lode        Not f'm figured
OperationState      UP
CAPllN' Path nru   1485
Link!U:liting      disabled
ApRole              RootAP
ApBac:khaul        802.11a
ApBac:khaulthannel 5805
ApBac:khaulSlot    1
ApBac:khaul1lgEnabled 0
ApBac:l<haul1xRate 24000
Ethernet Brl dglrg State 0
Public Safety State enabled
```

```
HJHAP111lehoi.I nesh adjacency ?
alI      HESH Adjacency Al I
child    HESH Adjacency Child
parent   MESH Adjacency Parent
oi
```

```
HJMap4#show mesh status ^
show MESH Status
MeshAP in state Maint
Uplink Backbone: Virtual-Dot11Radio0
Downlink Backbone: Dot11Radio1
Configured BGN: HuckJr
  rxNeighReq 129790 rxNeighRep 66976 txNeighReq 33938 txNeighRep 129790
  rxNeighReq 1147275 txNeighUpd 202060
  nextChan 0 nextant 0 downAnt 0 downChan 0 curAnts 0
  nextNeigh 1. malformedNeighPackets 4.poorNeighSnr 1
  blacklistPackets 0.insufficientMemory 0. authenticationFailures 0
  Parent Changes 3, Neighbor Timeouts 0
  Vector through 0017.94fe.c3bf:
    Vector ease 1 -1, FWD: 0017.94fe.c3bf
```

273949

```
HJMap4#show mesh forwarding link
Current mesh links:
-----
End Point   : 0017.94fe.c3bf
Adjacency   : Exists
Channel     : 161 on Dot11Radio1
Type        : 2
State       : 4
Bundle      : member
Bridge      : 1
swidb       : Virtual-Dot11Radio0
port state  : OPEN
```

273950



## Mesh Access Point Debug Commands

You can enter these commands directly on the mesh access point using the AP console port or you can use the remote debug feature from the controller.

- **debug mesh ethernet bridging**—Debugs Ethernet bridging.
- **debug mesh ethernet config**—Debugs access and trunk port configuration associated with VLAN tagging.
- **debug mesh ethernet registration**—Debugs the VLAN registration protocol. This command is associated with VLAN tagging.
- **debug mesh forwarding table**—Debugs the forwarding table containing bridge groups.
- **debugs mesh forwarding packet bridge-group**—Debugs the bridge group configuration.

## Defining Mesh Access Point Role

By default, AP1500s are shipped with a radio role set to MAP. You must reconfigure a mesh access point to act as a RAP.

## Backhaul Algorithm

A **backhaul** is used to create only the wireless connection between mesh access points.

The backhaul interface by default is 802.11a. You cannot change the backhaul interface to 802.11b/g.

The "auto" data rate is selected by default for AP1500s.

The backhaul algorithm has been designed to fight against stranded mesh access point conditions. This algorithm also adds a high-level of resiliency for each mesh node.

The algorithm can be summarized as follows:

- A MAP always sets the Ethernet port as the **primary backhaul** if it is UP; otherwise, it is the 802.11a radio (this feature gives the network administrator the ability to configure it as a RAP the first time and recover it in-house). For fast convergence of the network, we recommend that you do not connect any Ethernet device to the MAP for its initial joining to the mesh network.
- A MAP failing to connect to a WLAN controller on an Ethernet port that is UP, sets the 802.11a radio as the **primary backhaul**. Failing to find a neighbor or failing to connect to a WLAN controller via any neighbor on the 802.11a radio causes the **primary backhaul** to be UP on the Ethernet port again. A MAP gives preference to the parent which has the same BGN.
- A MAP connected to a controller over an Ethernet port does not build a mesh topology (unlike a RAP).
- A RAP always sets the Ethernet port as the **primary backhaul**.



---

**Note** Cisco Wave 2 APs operating as RAPs can fall back on Ethernet sooner than 15 minutes if the RAPs cannot find any valid uplink on the radio in 5 minutes' time. In such a case, the RAPs clear the blocked-listing on the wired port and try to fall back on the wired port.

---

- If the Ethernet port on a RAP is DOWN, or a RAP fails to connect to a controller on an Ethernet port that is UP, the 802.11a radio is set as the **primary backhaul**. Failing to find a neighbor or failing to connect to a controller via any neighbor on the 802.11a radio makes the RAP go to the SCAN state after 15 minutes and starts with the Ethernet port first.

Keeping the roles of mesh nodes distinct using the above algorithm greatly helps to avoid a mesh access point from being in an unknown state and becoming stranded in a live network.

## Passive Beacons (Anti-Stranding)

When enabled, passive beacons allows a stranded mesh access point to broadcast its debug messages over-the-air using a 802.11b/g radio. A neighboring mesh access point that is listening to the stranded mesh access point and has a connection to a controller, can pass those messages to the controller over CAPWAP. Passive beacons prevents a mesh access point that has no wired connection from being stranded.

Debug logs can also be sent as distress beacons on a nonbackhaul radio so that a neighboring mesh access point can be dedicated to listen for the beacons.

The following steps are automatically initiated at the controller when a mesh access point loses its connection to the controller:

- Identifies the MAC address of a stranded mesh access point
- Finds a nearby neighbor that is CAPWAP connected
- Sends commands through remote debug
- Cycles channels to follow the mesh access point

You only have to know the MAC address of the stranded AP to make use of this feature.

A mesh access point is considered stranded if it goes through a lonely timer reboot. When the lonely timer reboot is triggered, the mesh access point, which is now stranded, enables passive beacons, the anti-stranding feature.

This feature can be divided into three parts:

- Strand detection by stranded mesh access point
- Beacons sent out by stranded mesh access point
  - Latch the 802.11b radio to a channel (1,6,11)
  - Enable debugs
  - Broadcast the standard debug messages as distress beacons
  - Send Latest Crash info file
- Receive beacons (neighboring mesh access point with remote debugging enabled)

Deployed mesh access points constantly look for stranded mesh access points. Periodically, mesh access points send a list of stranded mesh access points and SNR information to the controller. The controller maintains a list of the stranded mesh access points within its network.

When the **debug mesh astools troubleshoot mac-addr start** command is entered, the controller runs through the list to find the MAC address of the stranded mesh access point.

A message is sent to the best neighbor to start listening to the stranded access point. The listening mesh access point gets the distress beacons from the stranded mesh access point and sends it to the controller.

Once a mesh access point takes the role of a listener, it does not purge the stranded mesh access point from its internal list until it stops listening to the stranded mesh access point. While a stranded mesh access point is being debugged, if a neighbor of that mesh access point reports a better SNR to the controller than the current listener by some percentage, then the listener of the stranded mesh access point is changed to the new listener (with better SNR) immediately.

End-user commands are as follows:

- **config mesh astools [enable | disable]**—Enables or disables the astools on the mesh access points. If disabled, APs no longer sends a stranded AP list to the controller.
- **show mesh astools stats**—Shows the list of stranded APs and their listeners if they have any.
- **debug mesh astools troubleshoot mac-addr start**—Sends a message to the best neighbor of the *mac-addr* to start listening.
- **debug mesh astools troubleshoot mac-addr stop**—Sends a message to the best neighbor of the *mac-addr* to stop listening.
- **clear mesh stranded [all | mac of b/g radio]**—Clears stranded AP entries.

The controller console is swamped with debug messages from stranded APs for 30 minutes.

## Misconfiguration of the Mesh Access Point IP Address

Although most Layer 3 networks are deployed using DHCP IP address management, some network administrators might prefer the manual IP address management and allocating IP addresses statically to each mesh node. Manual mesh access point IP address management can be a nightmare for large networks, but it might make sense in small to medium size networks (such as 10 to 100 mesh nodes) because the number of mesh nodes are relatively small compared to client hosts.

Statically configuring the IP address on a mesh node has the possibility of putting a MAP on a wrong network, such as a subnet or VLAN. This mistake could prevent a mesh access point from successfully resolving the IP gateway and failing to discover a WLAN controller. In such a scenario, the mesh access point falls back to its DHCP mechanism and automatically attempts to find a DHCP server and obtains an IP address from it. This fallback mechanism prevents a mesh node from being potentially stranded from a wrongly configured static IP address and allows it to obtain a correct address from a DHCP server on the network.

When you are manually allocating IP addresses, we recommend that you make IP addressing changes from the furthest mesh access point child first and then work your way back to the RAP. This recommendation also applies if you relocate equipment. For example, if you uninstall a mesh access point and redeploy it in another physical location of the mesh network that has a different addressed subnet.

Another option is to take a controller in Layer 2 mode with a RAP to the location with the misconfigured MAP. Set the bridge group name on the RAP to match the MAP that needs the configuration change. Add the MAP's MAC address to the controller. When the misconfigured MAP comes up in the mesh access point summary detail, configure it with an IP address.

## Misconfiguration of DHCP

Despite the DHCP fallback mechanism, there is still a possibility that a mesh access point can become stranded, if any of the following conditions exist:

- There is no DHCP server on the network.
- There is a DHCP server on the network, but it does not offer an IP address to the AP, or if it gives a wrong IP address to the AP (for example, on a wrong VLAN or subnet).

These conditions can strand a mesh access point that is configured with or without a wrong static IP address or with DHCP. Therefore, you must ensure that when a mesh access point is unable to connect after exhausting all DHCP discovery attempts or DHCP retry counts or IP gateway resolution retry counts, it attempts to find a controller in Layer 2 mode. In other words, a mesh access point attempts to discover a controller in Layer 3 mode first and in this mode, attempts with both static IP (if configured) or DHCP (if possible). The AP then attempts to discover a controller in Layer 2 mode. After finishing a number of Layer 3 and Layer 2 mode attempts, the mesh access point changes its parent node and re-attempts DHCP discovery. Additionally, the software exclusion-lists notes the parent node through which it was unable to obtain the correct IP address.

## Identifying the Node Exclusion Algorithm

Depending on the mesh network design, a node might find another node “best” according to its routing metric (even recursively true), yet it is unable to provide the node with a connection to the correct controller or correct network. It is the typical honeypot access point scenario caused by either misplacement, provisioning, design of the network, or by the dynamic nature of an RF environment exhibiting conditions that optimize the AWPP routing metric for a particular link in a persistent or transient manner. Such conditions are generally difficult to recover from in most networks and could blackhole or sinkhole a node completely, taking it out from the network. Possible symptoms include, but are not limited to the following:

- A node connects to the honeypot but cannot resolve the IP gateway when configured with the static IP address, or cannot obtain the correct IP address from the DHCP server, or cannot connect to a WLAN controller.
- A node ping-pongs between a few honeypots or circles between many honeypots (in worst-case scenarios).

Cisco mesh software resolves this difficult scenario by using a sophisticated node exclusion-listing algorithm. This node exclusion-listing algorithm uses an exponential backoff and advance technique much like the TCP sliding window or 802.11 MAC.

The basic idea relies on the following five steps:

**1. Honeypot detection**—The honeypots are first detected via the following steps:

A parent node is set by the AWPP module by:

- A static IP attempt in CAPWAP module.
- A DHCP attempt in the DHCP module.
- A CAPWAP attempt to find and connect to a controller fails.

**2. Honeypot conviction**—When a honeypot is detected, it is placed in a exclusion-list database with its conviction period to remain on the list. The default is 32 minutes. Other nodes are then attempted as parents in the following order, falling back to the next, upon failing the current mechanism:

- On the same channel.
- Across different channels (first with its own bridgegroupname and then with default).
- Another cycle, by clearing conviction of all current exclusion-list entries.

- Rebooting the AP.
3. Nonhoneypot credit—It is often possible that a node is not really a honeypot, but appears to be due to some transient back-end condition, such as the following:
    - The DHCP server is either not up-and-running yet, has failed temporarily, or requires a reboot.
    - The WLAN controller is either not up-and-running yet, has failed temporarily, or requires a reboot.
    - The Ethernet cable on the RAP was accidentally disconnected.

Such nonhoneypots must be credited properly from their serving times so that a node can come back to them as soon as possible.
  4. Honeypot expiration—Upon expiration, an exclusion-list node must be removed from the exclusion-list database and return to a normal state for future consideration by AWPP.
  5. Honeypot reporting—Honeypots are reported to the controller via an LWAPP mesh neighbor message to the controller, which shows these on the Bridging Information page. A message is also displayed the first-time an exclusion-listed neighbor is seen. In a subsequent software release, an SNMP trap is generated on the controller for this condition so that Cisco Prime Infrastructure can record the occurrence.

**Figure 5: Excluded Neighbor**

All APs > sjc10-p1012-map1:62:40:d0 > Bridging Details < Back

Bridging Details		Bridging Links	
AP Role	MeshAP	Mesh Type	AP Name/Radio M
Bridge Group Name	betamesh	Parent	sjc14-41a-rap3-5e:9
Backhaul Interface	802.11a	<b>Excluded Neighbor</b>	00:0B:85:53:4B:30
Switch Physical Port	29	Neighbor	00:0B:85:5C:B8:A0
Routing State	Maintenance	Neighbor	00:0B:85:5C:B9:80
Malformed Neighbor Packets	0	Neighbor	00:0B:85:5F:FA:50
Poor Neighbor SNR reporting	1	Neighbor	00:0B:85:5F:FE:E0
Blacklisted Packets	212	Neighbor	00:0B:85:5F:FF:40
Insufficient Memory reporting	0	Neighbor	00:0B:85:5F:FF:E0

Because many nodes might be attempting to join or rejoin the network after an expected or unexpected event, a hold-off time of 16 minutes is implemented, which means that no nodes are exclusion-listed during this period of time after system initialization.

This exponential backoff and advance algorithm is unique and has the following properties:

- It allows a node to correctly identify the parent nodes whether it is a true honeypot or is just experiencing temporary outage conditions.
- It credits the good parent nodes according to the time it has enabled a node to stay connected with the network. The crediting requires less and less time to bring the exclusion-list conviction period to be very low for real transient conditions and not so low for transient to moderate outages.
- It has a built-in hysteresis for encountering the initial condition issue where many nodes try to discover each other only to find that those nodes are not really meant to be in the same network.
- It has a built-in memory for nodes that can appear as neighbors sporadically so they are not accidentally considered as parents if they were, or are supposed to be, on the exclusion-list database.

The node exclusion-listing algorithm guards the mesh network against serious stranding. It integrates into AWPP in such a way that a node can quickly reconverge and find the correct network.

## Throughput Analysis

Throughput depends on packet error rate and hop count.

Capacity and throughput are orthogonal concepts. Throughput is one user's experience at node N and the total area capacity is calculated over the entire sector of N-nodes and is based on the number of ingress and egress RAP, assuming separate noninterfering channels.

For example, 4 RAPs at 10 Mbps each deliver 40 Mbps total capacity. So, one user at 2 hops out, logically under each RAP, could get 5 Mbps each of TPUT, but consume 40 Mbps of the backhaul capacity.

With the Cisco Mesh solution, the per-hop latency is less than 10 msec, and the typical latency numbers per hop range from 1 to 3 msec. Overall jitter is also less than 3 msec.

Throughput depends on the type of traffic being passed through the network: User Datagram Protocol (UDP) or Transmission Control Protocol (TCP). UDP sends a packet over Ethernet with a source and destination address and a UDP protocol header. It does not expect an acknowledgement (ACK). There is no assurance that the packet is delivered at the application layer.

TCP is similar to UDP but it is a reliable packet delivery mechanism. There are packet acknowledgments and a sliding window technique is used to allow the sender to transmit multiple packets before waiting for an ACK. There is a maximum amount of data the client transmits (called a TCP socket buffer window) before it stops sending data. Sequence numbers track packets sent and ensure that they arrive in the correct order. TCP uses cumulative ACKs and the receiver reports how much of the current stream has been received. An ACK might cover any number of packets, up to the TCP window size.

TCP uses slow start and multiplicative decrease to respond to network congestion or packet loss. When a packet is lost, the TCP window is cut in half and the back-off retransmission timer is increased exponentially. Wireless is subject to packet loss due to interference issues and TCP reacts to this packet loss. A slow start recovery algorithm is also used to avoid swamping a connection when recovering from packet loss. The effect of these algorithms in a lossy network environment is to lessen the overall throughput of a traffic stream.

By default, the maximum segment size (MSS) of TCP is 1460 bytes, which results in a 1500-byte IP datagram. TCP fragments any data packet that is larger than 1460 bytes, which can cause at least a 30-percent throughput drop.

**Wireless**

All APs > Details for AP1562.EC4A.4D30

**General** | Credentials | Interfaces | High Availability | Inventory | Mesh | Advanced

**General**

AP Name: AP1562.EC4A.4D30  
 Location: default location  
 AP MAC Address: 00:62:ec:4a:4d:30  
 Base Radio MAC: 00:62:ec:06:5d:40  
 Admin Status: Disable  
 AP Mode: **Bridge** (selected)  
 AP Sub Mode: local  
 Operational Status: monitor  
 Port Number: Sniffer  
 Venue Group: SE-Connect  
 Venue Type: Unspecified

**Versions**

Primary Software Version	8.5.103.0
Backup Software Version	8.5.1.204
Predownload Status	None
Predownloaded Version	None
Predownload Next Retry Time	NA
Predownload Retry Count	NA
Boot Version	1.1.2.4
IOS Version	8.5.103.0
Mini IOS Version	0.0.0.0

**IP Config**

CAPWAP Preferred Mode	Ipv4 (Global Conf)
DHCP Ipv4 Address	10.70.0.129
Static IP (Ipv4/Ipv6)	<input type="checkbox"/>

**Fabric**

Fabric Status	Disabled
Fabric L2 Instance ID	0

**Network Spectrum**

Interface Key	9118035629CC881ADEFEE36B765B885A
---------------	----------------------------------

**GPS Location**

GPS Present	No
-------------	----

