Troubleshoot MAC Address Flap Notification Error

Contents

MAC Address Flap Notification

ICSeverity

5 - Notice

Impact

These messages can be investigated to ensure that a forwarding loop does not exist.

Description

This notification message is generated by the switch when it detects a MAC address flapping event on the network. A MAC address flapping event is detected when a switch receives packets from the same Source MAC address into two different interfaces. Cisco Catalyst switches notify when the same MAC address is detected on multiple switch ports, causing the switch to constantly change the port associated with the MAC address, and alert via this syslog that contains the MAC address of the host, VLAN, and ports between which the MAC address is flapping. Given that this behavior can be caused due to multiple reasons, identifying the underlying cause of MAC address flapping is important to ensure the stability and performance of the network.

SyslogMessage

```
SW_MATM-4-MACFLAP_NOTIF
```

MessageSample

```
Apr 26 12:27:55 <> %SW_MATM-4-MACFLAP_NOTIF: Host mac address in vlan X is flapping between port PoX an
```

ProductFamily

- Cisco Catalyst 9300 Series Switches
- Cisco Catalyst 9400 Series Switches
- Cisco Catalyst 9200 Series Switches
- Cisco Catalyst 9500 Series Switches
- Cisco Catalyst 9600 Series Switches
- Cisco Catalyst 3850 Series Switches
- Cisco Catalyst 3650 Series Switches
- Cisco Catalyst 6000 Series Switches
- Cisco Catalyst 6800 Series Switches
- Cisco Catalyst 4500 Series Switches
- Cisco Catalyst 4900 Series Switches
- Cisco Catalyst 3750-X Series Switches
- Cisco Catalyst 3850-X Series Switches
- Cisco Catalyst 2960 Series Switches

Regex

N/A

Recommendation

There are many possible causes for this error, some of which can indicate a serious network problem. The 3 most common are explained in detail below:

1. Wireless client movement (no network impact).

2. Virtual address movement from redundant systems or duplicated Virtual Machines (moderate network impact).

3. Layer-2 loops (high network impact)

#1 Details: Wireless client movement is often expected, and can usually be safely ignored assuming there are no service impacts observed. Clients roaming between APs that are not using CAPWAP back to a wireless controller, or roam between APs controlled by two different wireless controllers, are likely to generate this log. The time between logs generated for the same mac address can be several seconds or several minutes apart. If you see a single mac address is moving multiple times per second, that can indicate a more serious problem and additional troubleshooting can be required.

#2 Details: Some redundant systems or devices operating in an active/standby state can share a common virtual IP and mac address, with only the active device using it at any point in time. If both devices unexpectedly become active and both begin using the virtual address, this error can be seen. Using a combination of the interfaces mentioned in the log and the **show mac address-table address vlan** command trace the path of this mac through the network to determine where and which devices are generating traffic from the shared mac. Depending on the nature of the devices generating the moves, additional troubleshooting of their redundancy states can be required. #3 Details: L2 loops often generate a large number of mac move errors in a very short period of time (at least one per second, often more). Logs can typically be for a single, or a small number of mac addresses, and users can experience an impact to the network. Routing and layer 2 protocols can often fail resulting in additional logs and general instability

being created. To troubleshoot an L2 loop, run the command **show int | in is up|input rate** and note all of the active interfaces that show an extremely high volume of input packets per second (generally speaking, this can be a very large 6, 7, or 8+ digit number depending on the speed of the interface). There is likely to be only 1 or 2 interfaces with an abnormally high input rate. Do not focus on output rates, and do not focus on spanning-tree TCNs. Once the high input interface is identified, use CDP, LLDP, or your interface descriptions/network diagram to log into the neighboring device connected to that port, and run the **show int | in is up|input rate** command again and repeat the process of tracing the interfaces with abnormal input rates. Keep track of the interfaces and hostnames as you trace them through the network. Continue checking neighbors and looking at input rates until you run out of input ports, and you run out of neighbors or end up back on the device you already checked. One of two possible outcomes can happen during this methodology: If you end up with a port that has no CDP, LLDP, or known neighbor, but a very high input rate, administratively shut it down. This interface is likely the ultimate source, or is a contributor to the loop. Wait 60 seconds for the network to stabilize, and if a loop condition is still seen, keep the interface shutdown and start the process over, since its possible there is a 2nd source on the network. If you end up on a device you already checked, this indicates the loop prevention protocol in use (Spanning-tree being the most common) has failed somewhere. For spanning-tree networks, identify which switch in the path you traced is expected to be root, and work backwards from that device to determine which interface can be in a blocking state within your traced path. Once the interface that can be blocking (but is in forwarding state) is found, administratively shut it down. Wait 60 seconds and check the network for stability. If the loop persists, keep the interface shutdown and repeat this process.

Commands

#show version

#show logging

#show spanning-tree

#show mac-address-table

#show mac address-table