# Support VMware vSphere vMotion on Catalyst 9800

## Contents


**Introduction**

**Prerequisites**

- Requirements
- Topology

**Test results**

**Summary**


## Introduction

This document describes the tests done to verify vMotion support for the C9800-CL that runs on vSphere ESXi.

## Prerequisites

C9800-CL is the virtual machine form factor of the Catalyst 9800 Wireless LAN Controller. You can use VMware vSphere vMotion to perform a zero-downtime live migration of Catalyst 9800-CL from one host server to another. This capability is possible across vSwitches and clusters. The goal is that, during the C9800-CL live migration, the wireless network remains up and wireless users continue to have the connectivity they need.

vMotion can be done manually or as part of a VMware vSphere Distributed Resource Scheduler (DRS) configuration. DRS spreads the virtual machine workloads across vSphere hosts inside a cluster and monitors available resources for you. Based on your automation level, DRS migrates virtual machines to other hosts within the cluster to maximize performance.

VMware functionalities, such as, vMotion, Snapshot, Distributed Resource Scheduler (DRS), vNIC Teaming and SR-IOV modes are supported. However, cloning from snapshots is not supported.

Also, vMotion, DRS, Snapshots, and vNIC Teaming are not supported when SR-IOV mode is enabled.

For more information, see the Cisco Catalyst 9800-CL Wireless Controller for Cloud Data Sheet.

### Requirements

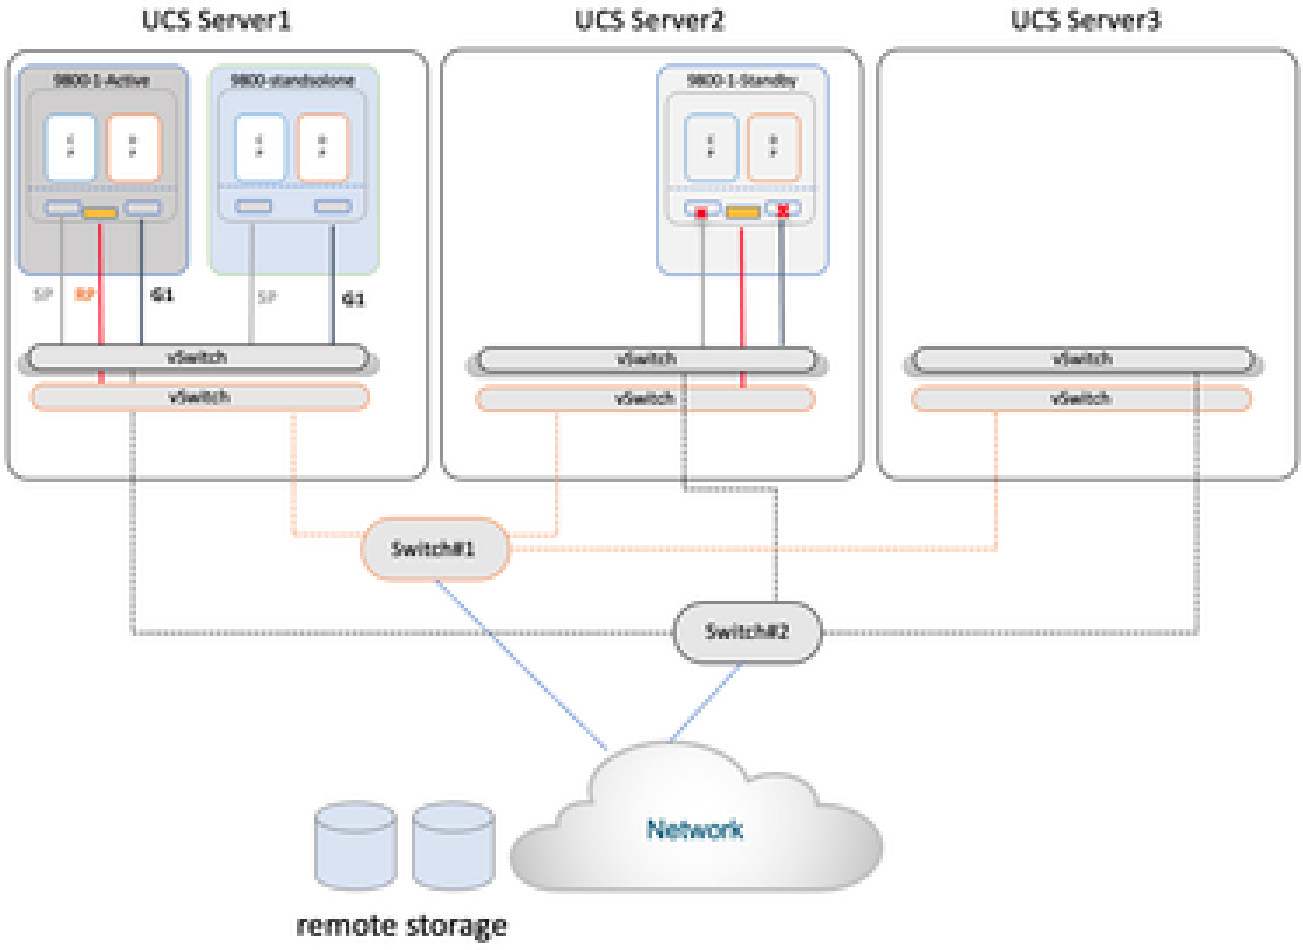
- Use recommended tested software releases:
  - ESXi vCenter 6.7 or later
  - C9800-CL software: 17.9.2 and later
- Latency (RTT) between the remote storage to the server where C9800-CL runs must be < 60 ms
- C9800-CL VM must not have any ESXi host specific correspondence like CD/DVD, serial console port connection, and so on.
- Configure vMotion as per the VMware guidelines for host, remote shared storage and networking here
- Comply with VMware network requirements for vMotion here .

### Topology

For these verification tests a simple topology was used with three different server hosts and iSCSI remote storage (NFS storage can be used as well). The remote storage leverages 10 Gbps connection to the servers. On the ESXi host, one C9800-CL VM is created in standalone mode and two other C9800-CL virtual machines configured for Stateful Switchover High Availability (SSO HA). The HA pair is created across two different servers for physically redundancy and to be able to migrate both active and standby WLC separately. Each C9800-CL VM is connected to the virtual switch by the usage of three ports:

- G1 > SP port (optional)
- G2 > Trunk port for Wireless Management Interface (WMI) VLAN and client VLANs if present
- G3 > RP port. This is for the SSO cluster creation. Not connected for the standalone mode

Each host server has a dedicated physical port and dedicated switch (switch#1) to connect the RP ports together through a L2 link, across the servers. The other two physical ports are connected to a separated uplink switch (switch#2). A diagram that represents the test topology:




# Test results

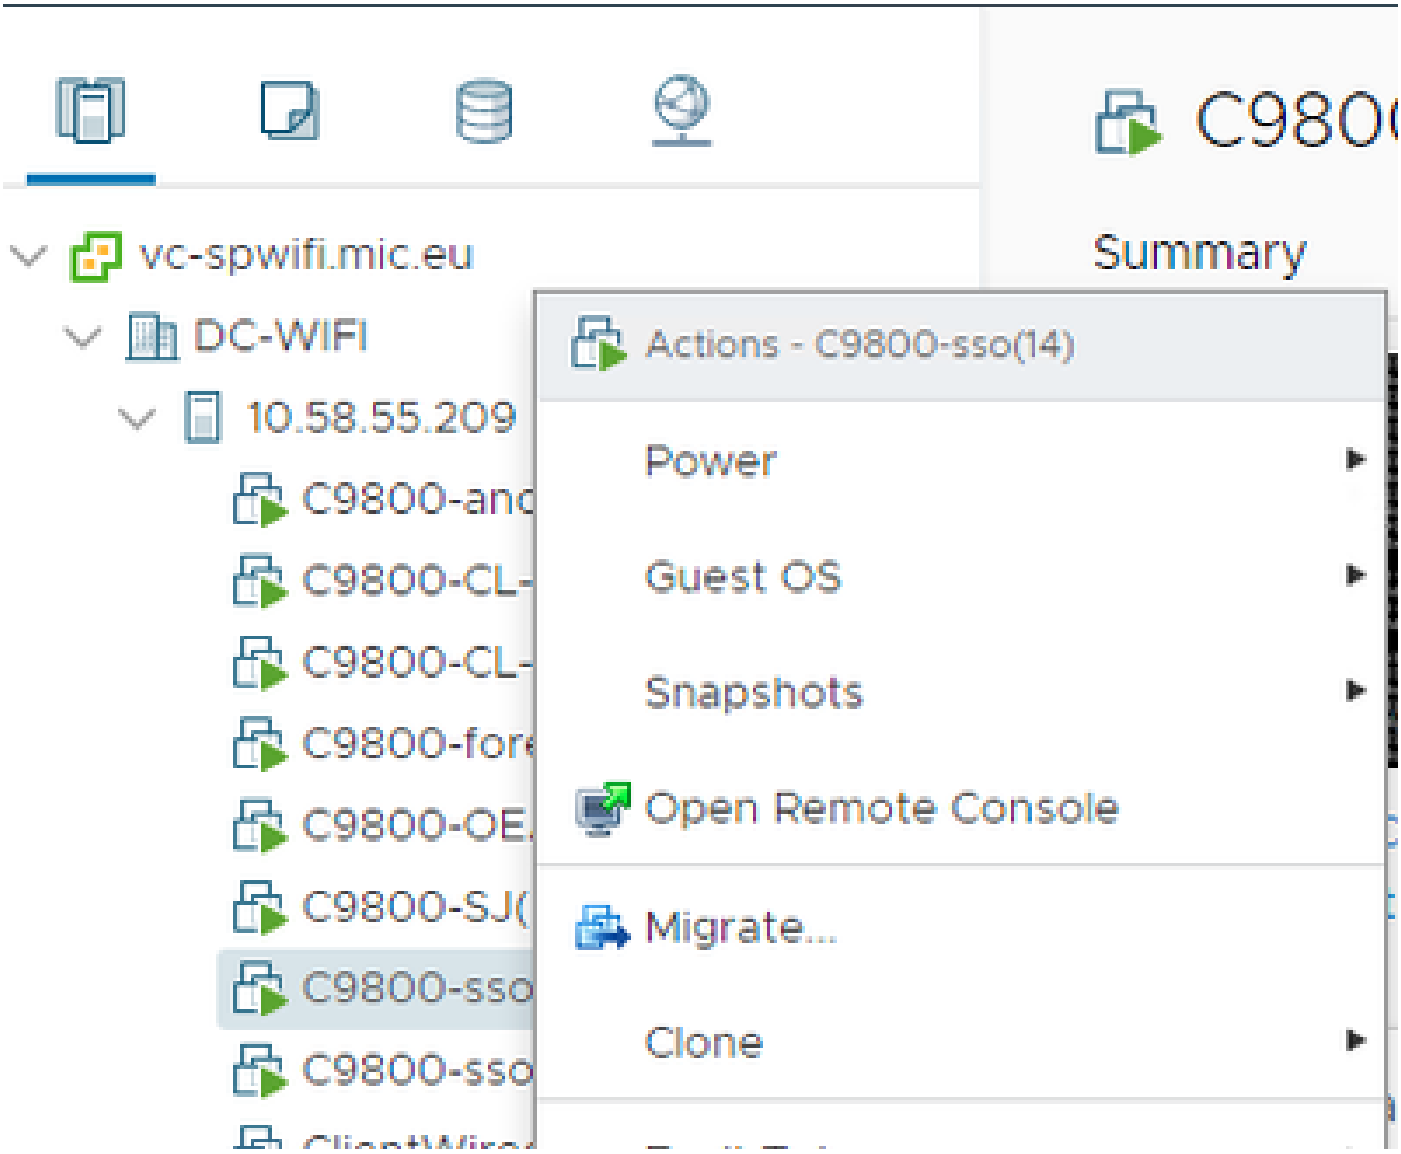
For these tests, two migration scenarios where considered:

1. A standalone C9800-CL is migrated between server #1 and server #2
2. A pair of C9800-CL configured as in SSO high availability. In this case first the active is migrated between server #1 and server #3 and then the standby WLC is migrated from server #2 to server #3

In both cases, all the three different types of vMotion migration were tested: compute resource only, storage

only, both compute and storage.

To trigger vMotion, just right click on the VM and click migrate:




Select the type of migration and go through the steps:




Here is the result of each test:

| Test | Standalone | vMotion type | Observations/Comments |
|---|---|---|---|

| | | | |
|---|---|---|---|
| | **C9800-CL** | | |
| 1 | | Compute resource only | **Not Supported**: APs and clients drop seen, which recover after some time, due to Virtual Guest Tagging (802.1q VLAN) issue: KB article<br>**Workaround**: Start continuous ping from the controller to any wired network device |
| 2 | | Storage only | **Supported**: APs and Clients are stable, single ping drop is seen |
| 3 | | Compute resource and storage | **Not Supported:** APs and clients drop seen, which recover after some time, due to Virtual Guest Tagging (802.1q VLAN) issue: KB article<br>**Workaround**: Start continuous ping from the controller to any wired network device |
| **Test** | **SSO Active**<br>**HA keepalive: 100ms** | **vMotion type** | |
| 4 | | Compute resource only | **Supported:** Traffic is stable on active, standby stack merge reload seen due to HA RP keepalives expired |
| 5 | | Storage only | **Supported:** Traffic is stable, most of the time RP comes up before RP keepalives timer expired so no stack merge is seen |
| 6 | | Compute resource and storage | **Supported:** Standby went to standby recovery state and reload due to stack merge. |
| **Test** | **SSO Active**<br>**HA keepalive: 200ms** | **vMotion type** | |
| 7 | | Compute resource only | **Supported:** APs and Clients are stable, single ping drop is seen on active, standby also stable |
| 8 | | Storage only | **Supported:** APs and Clients are stable, single ping drop is seen on active, stand also stable |
| 9 | | Compute resource and | **Supported:** APs and Clients are stable, single ping drop is seen on active, stand also stable |

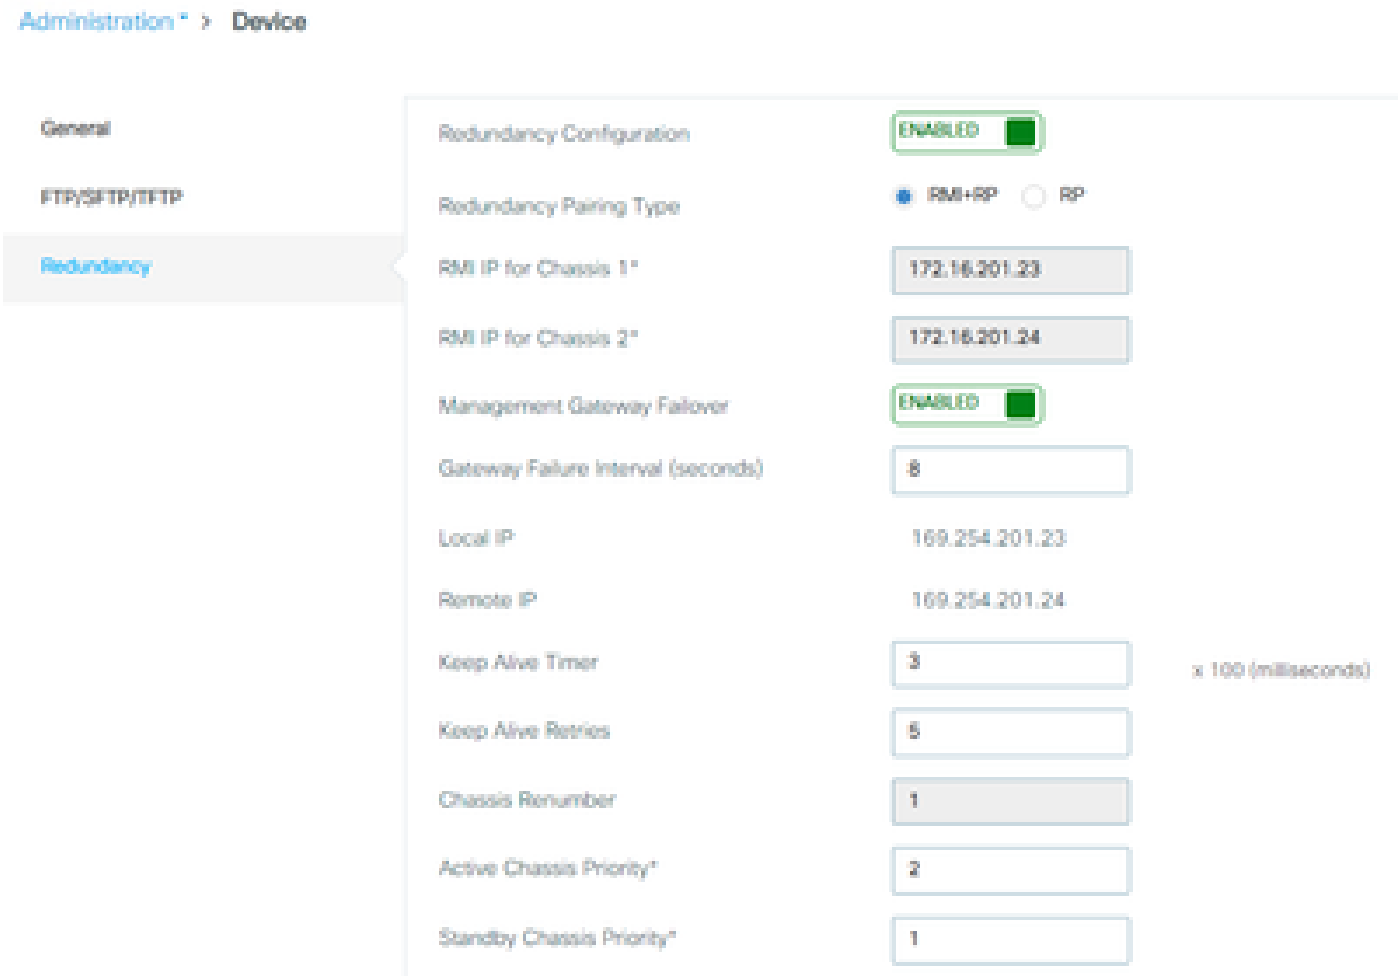
| | | | |
|---|---|---|---|
| | | storage | |
| **Test** | **SSO Standby**<br>**HA keepalive-100ms** | **vMotion type** | |
| 10 | | Compute resource only | **Supported:** APs and Clients are stable on active and stand also stable after the vMotion operation; sometimes standby stack merge reloads seen. |
| 11 | | Storage only | **Supported:** APs and Clients are stable on active and stand also stable after the vMotion operation; sometimes standby stack merge reloads seen. |
| 12 | | Compute resource and storage | **Supported:** APs and Clients are stable on active and stand also stable after the vMotion operation; sometimes standby stack merge reloads seen. |
| **Test** | **HA Standby**<br>**HA keepalive-200ms** | | |
| 13 | | Compute resource only | **Supported:** APs and Clients are stable on active and stand also stable after the vMotion operation |
| 14 | | Storage only | **Supported:** APs and Clients are stable on active and stand also stable after the vMotion operation |
| 15 | | Compute resource and storage | **Supported:** APs and Clients are stable on active and stand also stable after the vMotion operation |

As seen in this table, vMotion fails in the first and third scenario (test #1 and #3) with standalone mode C9800-CL, as it performs a compute or compute and storage migration; in this case the MAC and IP address of the C9800-CL's WMI moves to the new host and hence to a different switch port. vMotion is unable to send a Reverse Address Resolution Protocol (RARP) for the C9800-CL wireless management VLAN as the ESXi host cannot identify which VLAN is in use by the guest operating system that runs in the virtual machine. To support this scenario, you need to implement a workaround: start a continuous ping from the C9800-CL to any wired host before it performs the migration; this triggers the switch network to learn about the new location (port) of the VM and hence converge faster.

In the analogue migration case with HA SSO (test #4, for example), the Redundancy Management Interface (RMI) is leveraged to check reachability to the gateway and between Active and Standby, and hence it generates the traffic that keep the MAC address table on the switch updated and the problem doesn't

happen.

**Recommendation**: For best results, it is recommended to configure RP port keepalives to at least twice the default 100 ms keepalive (set it to 200 ms). If the network between storage and hosts can become busy and increase latency, consider to set the keepalives timer to 300 ms. To configure the keepalive timer on the GUI, go to Administration > Device > Redundancy:




On the CLI use this command in exec mode (not in configuration mode!)

```
C9800-SSO#chassis redundancy keep-alive timer 3
```

To verify, use this show command:

```
C9800-SSO#sh chassis ha-status active

                    My state = ACTIVE
                  Peer state = STANDBY HOT
        Last switchover reason = none
```

```
          Last switchover time = none

                 Image Version = 17.9.1


Chassis-HA   Local-IP         Remote-IP        MASK             HA-Interface

-------------------------------------------------------------------------------

This Boot: 169.254.201.23   169.254.201.24   255.255.255.0


Next Boot: 169.254.201.23   169.254.201.24   255.255.255.0



Chassis-HA   Chassis#    Priority        IFMac Address         Peer-timeout(ms)*Max-retry

Shape-------------------------------------------------------------------------------------

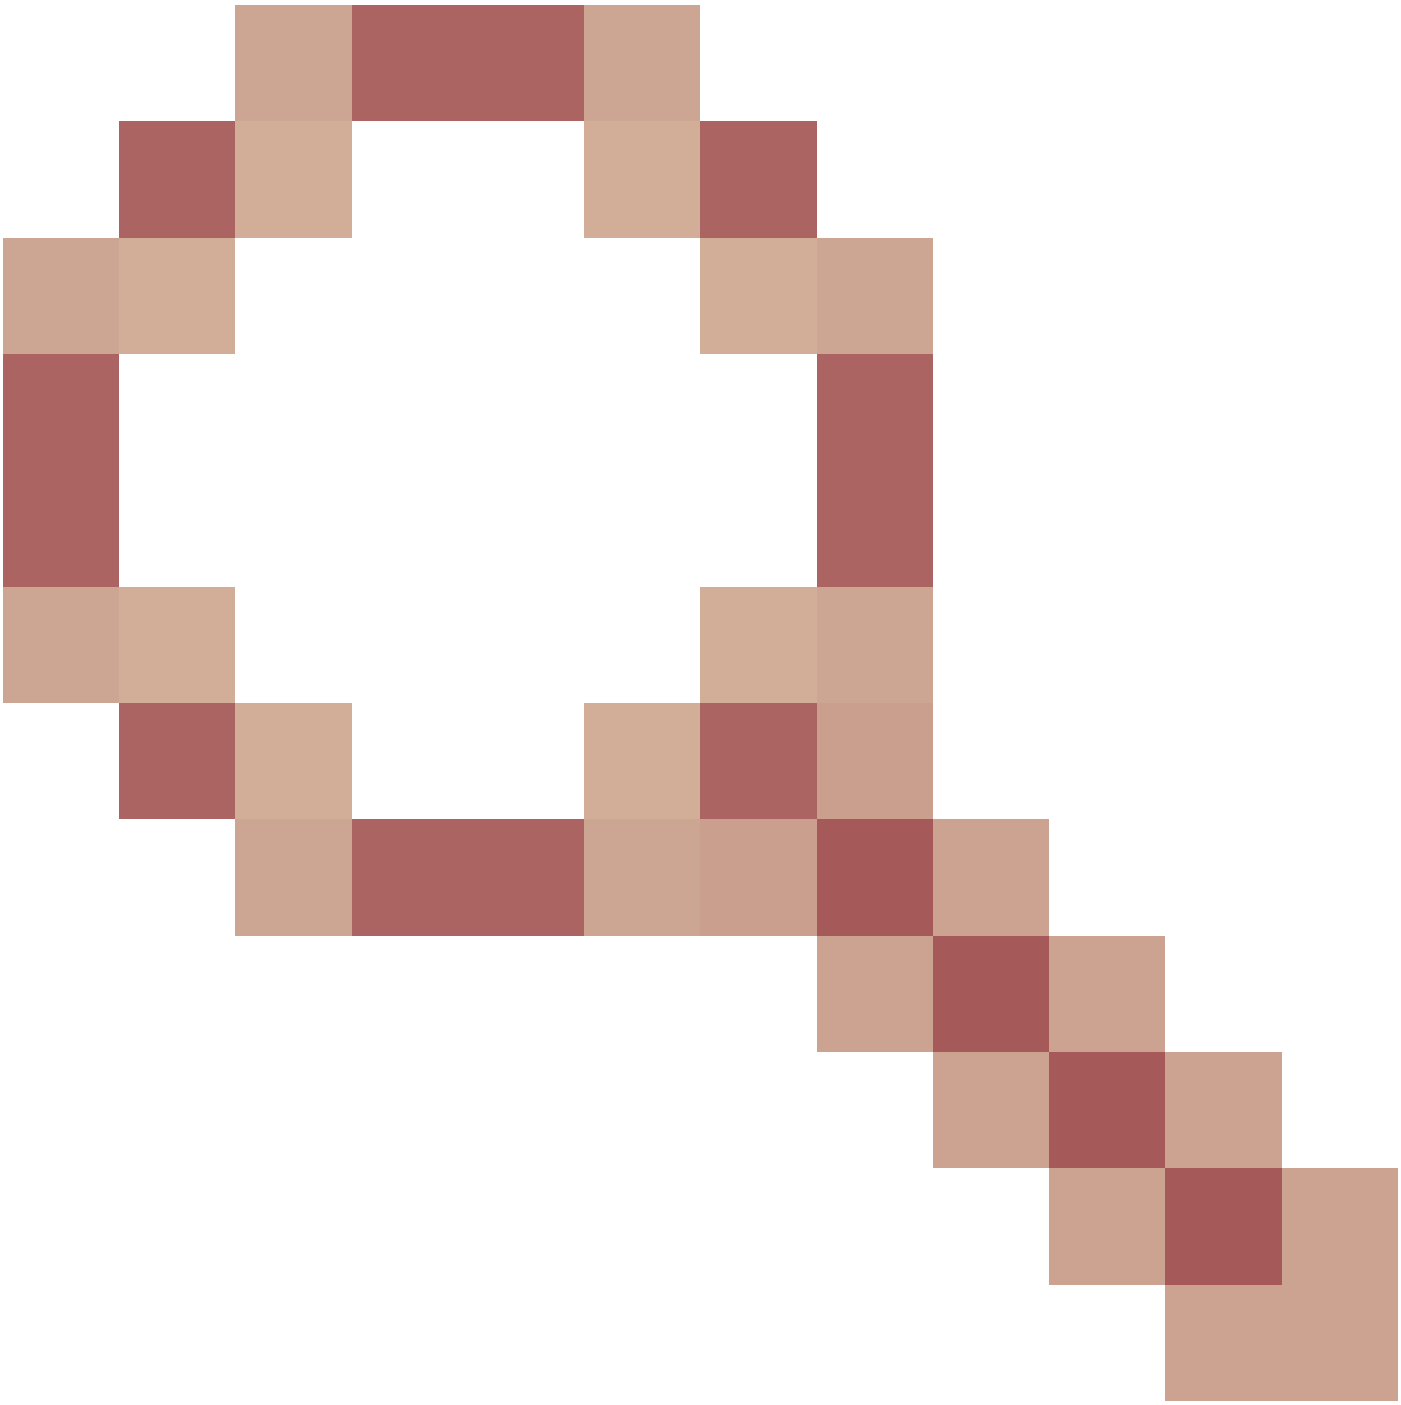
This Boot:    1             1                                     300*5


Next Boot:    1             1                                     300*5
```

Resolved caveats:

These are the caveats fixed in 17.9.2:

Cisco bug ID CSCwd17349



- C9800: Active chassis can get stuck during the SSO failover on 17.9

## Summary

VMware vSphere vMotion can be leveraged to migrate the C9800-CL VM from one host to the other without impact to the wireless network operations.  vMotion is officially supported on the C9800-CL as from release 17.9.2.