

Cisco Nexus Dashboard on KVM

Host Maintenance, Update Validation, and Recovery Guidance

August 11, 2026



Contents

- Purpose 3
- Audience 3
- Scope 3
- Responsibility Model 3
- Possible Effects of Host-Level Updates 3
- Availability Principles 4
- Before Host Maintenance 4
- Recommended Maintenance Sequence 5
- Post-Update Validation 5
- Recovery Guidance 6
- Information to Preserve for Troubleshooting 6
- Platform Position 7
- Recommendations 7

Purpose

Provide platform-level guidance for maintaining Nexus Dashboard availability when users patch, upgrade, or otherwise modify the Linux KVM host. This document intentionally avoids Linux administration procedures and product-specific host commands.

Audience

Nexus Dashboard administrators, virtualization administrators, and operations teams.

Scope

Cisco Nexus Dashboard virtual machines can be deployed on supported KVM environments running qualified Linux distributions and releases. Users might need to update the host kernel, virtualization packages, device drivers, firmware, or related host components to address security vulnerabilities, defects, or lifecycle requirements.

These host-level changes can affect VM startup, device presentation, storage access, networking, timing, performance, or the ability of a Nexus Dashboard node to rejoin the cluster. The objective of this guidance is to reduce operational risk, preserve cluster availability, and define the correct validation and recovery sequence.

Important distinction

A failure following a host update does not, by itself, indicate a Nexus Dashboard software defect or a Linux distribution qualification issue. Qualification confirms support for the documented distribution and release; it does not guarantee that every future host patch, kernel build, driver, firmware, or user-specific virtualization configuration will preserve all VM mappings and behavior unchanged.

Responsibility Model

Owner	Primary responsibility
User / KVM host administrator	Plan and execute host maintenance; preserve VM definitions and infrastructure configuration; validate compute, storage, networking, device presentation, and virtualization services; restore the VM environment if host-level changes affect it.
Nexus Dashboard administrator	Confirm cluster health before maintenance; coordinate one-node-at-a-time operations; validate node and cluster health after each host is returned to service; perform documented Nexus Dashboard recovery or node replacement procedures when required; and avoid concurrent disruptive work.
Linux, virtualization, or hardware vendor	Investigate host regressions, driver or firmware behavior, device enumeration changes, virtualization service failures, and other infrastructure problems introduced or exposed by host maintenance.

Possible Effects of Host-Level Updates

Host maintenance can introduce changes that prevent a VM from starting or can allow the VM to start while leaving the Nexus Dashboard node partially functional. Examples include:

-
- Changes in physical or virtual device discovery, numbering, addressing, or attachment.
 - Loss or modification of VM storage paths, storage pools, logical volumes, mounts, permissions, or device ordering.
 - Changes to Linux bridges, bonds, VLANs, virtual switches, interface naming, MAC presentation, MTU, routing, or firewall behavior.
 - Changes in KVM, QEMU, libvirt, kernel modules, drivers, IOMMU, SR-IOV, or passthrough behavior.
 - Changes in CPU feature exposure, machine type, memory allocation, timekeeping, or virtualization performance.
 - VM autostart failures, stale or altered persistent VM definitions, or changes in service startup ordering.
 - Unexpected interaction with firmware, BIOS, NIC, HBA, RAID, or platform updates performed with the host maintenance.
 - A host that appears healthy while one or more VM dependencies are unavailable or mapped differently.

Example

A kernel or firmware update might change PCI device enumeration. If a VM definition references the previous address, the VM might fail to start. PCI address change is one example of a broader class of host configuration and device-presentation changes.

Availability Principles

- Perform maintenance on only one KVM host and one Nexus Dashboard node at a time.
- Confirm that the Nexus Dashboard cluster is healthy before starting each maintenance action.
- Do not proceed to the next host until the current node has returned and the cluster is fully healthy.
- Avoid combining host patching with Nexus Dashboard upgrade, node replacement, restore, or other major cluster operations.
- Maintain a tested rollback path for the host and a recoverable copy of the VM definition and host configuration.
- Treat the VM, host networking, host storage, and any device mappings as a single operational dependency set.
- Use the Nexus Dashboard node replacement workflow only after host and VM recovery options have been exhausted.

Before Host Maintenance

Validate Nexus Dashboard Readiness

- All Nexus Dashboard nodes are reachable and healthy.
- The cluster has no unresolved critical faults that could reduce availability during maintenance.
- No upgrade, backup, restore, application installation, node replacement, or other disruptive operation is in progress.

- Applications and services are in an acceptable state for one node to be temporarily unavailable.
- A current Nexus Dashboard backup exists where required by the deployment and operational policy.

Preserve the Host and VM Baseline

The host administrator should preserve enough information to reconstruct the pre-maintenance VM environment. At minimum, retain:

- The persistent VM definition, VM identity, assigned CPU and memory, machine type, and autostart behavior.
- Virtual disk identities, attachment order, target devices, backing storage, and storage access requirements.
- Virtual NIC identities, MAC addresses, network attachments, bridge or virtual-switch mappings, VLANs, and MTU expectations.
- Any host-device, SR-IOV, passthrough, accelerator, or other direct device assignments.
- The current host OS, kernel, virtualization stack, driver, and firmware baseline.
- A supported host rollback method in case the updated environment cannot restore the VM correctly.

Establish Maintenance Ownership

Before maintenance begins, identify the Nexus Dashboard administrator, KVM host administrator, and escalation owner. The host administrator should be available until the node has successfully returned to the cluster, not only until the host reboot completes.

Recommended Maintenance Sequence

- Step 1.** Confirm that the Nexus Dashboard cluster is healthy and suitable for single-node maintenance.
- Step 2.** Select one Nexus Dashboard node and its corresponding KVM host.
- Step 3.** Gracefully shut down or place the selected node into the documented maintenance state.
- Step 4.** Confirm that the remaining cluster continues to operate as expected.
- Step 5.** Preserve the current host and VM baseline and verify that a rollback path is available.
- Step 6.** Apply the required host updates and reboot or restart host services as required by the host maintenance plan.
- Step 7.** Validate the KVM host and all VM dependencies before starting the Nexus Dashboard VM.
- Step 8.** Start the VM and validate its console, virtual hardware, storage, and network connectivity.
- Step 9.** Confirm that the Nexus Dashboard node rejoins the cluster and that cluster health is fully restored.
- Step 10.** Proceed to the next host only after the current node and cluster are stable.

Post-Update Validation

Host and VM Validation

- The virtualization services are operational and the persistent VM definition is intact.
- All required virtual disks are present, accessible, and attached in the expected order.
- All required virtual NICs are present with the expected identities and network attachments.

- Host bridges, bonds, VLANs, storage, and other dependencies required by the VM are operational.
- Any direct device assignments or passthrough mappings still identify the intended devices.
- The VM starts without host-level errors and reaches its normal operating state.
- CPU, memory, timing, and performance are consistent with the supported deployment requirements.

Nexus Dashboard Validation

- The node is reachable on the expected management and data networks.
- The node retains its expected identity, addresses, and cluster membership.
- The node rejoins the cluster without requiring manual re-creation.
- Cluster quorum and service health are restored.
- Applications and platform services return to their expected state.
- No new critical faults are present that are attributable to the maintenance event.

Recovery Guidance

Recovery rule

Recover the KVM host and VM environment first. Nexus Dashboard recovery starts only after the VM can boot with the expected storage, networking, and virtual hardware.

VM Does Not Start

The host administrator should compare the post-update environment with the preserved baseline and correct any missing or changed VM dependency. Typical areas include device mappings, storage availability, network attachments, virtualization services, permissions, drivers, and persistent VM configuration. If the updated host cannot reproduce the required environment, use the host rollback plan or engage the Linux, virtualization, or hardware vendor.

VM Starts but the Node Is Unreachable

Validate that all expected virtual NICs, MAC addresses, network attachments, VLANs, MTU settings, routing, and host filtering behavior remain intact. Also verify that the VM has access to the expected boot and data storage and that device ordering has not changed.

VM Starts but the Node Does Not Rejoin the Cluster

After host and VM dependencies are confirmed, the Nexus Dashboard administrator should validate node identity, cluster connectivity, time synchronization, storage health, and service state, and perform the documented recovery procedure if required.

Node Cannot Be Recovered

If the KVM environment is healthy but the existing Nexus Dashboard node cannot be returned to service, follow the documented Nexus Dashboard node replacement procedure. Node replacement should not be used to bypass an unresolved host problem, because the replacement VM can encounter the same underlying condition.

Information to Preserve for Troubleshooting

For efficient troubleshooting and recovery, preserve the following before and after the host update:

- Host OS, kernel, virtualization stack, driver, and firmware versions.
- Maintenance timeline and exact components changed.
- Pre-update and post-update VM definitions.
- Host, virtualization, storage, and network error messages associated with the failed VM or node.
- Identification of any device, disk, interface, address, mapping, or service that changed.
- Nexus Dashboard cluster health before maintenance and after the node was restarted.
- Whether rollback restores the previous behavior.

Platform Position

Cisco validates Nexus Dashboard on documented Linux distributions and supported OS releases. Users might apply host updates required for security, defect correction, or lifecycle management. The user is responsible for ensuring that the updated KVM host continues to provide the VM configuration, storage, networking, device mappings, and virtualization behavior required by Nexus Dashboard. If a host update changes or disrupts those dependencies, the KVM host administrator must restore the VM environment. After the VM is operational, the Nexus Dashboard administrator must validate node and cluster health and perform the documented recovery or node replacement procedure when required.

Recommendations

- Use a rolling, one-host-at-a-time maintenance model.
- Validate cluster health before and after every host maintenance action.
- Preserve the VM and host baseline before applying updates.
- Maintain a tested rollback path for the host.
- Do not assume that a successful host boot means all VM dependencies are intact.
- Validate VM storage, networking, device presentation, and virtual hardware before evaluating Nexus Dashboard.
- Engage the appropriate Linux, virtualization, or hardware vendor for host-level regressions or configuration changes.
- The Nexus Dashboard administrator should complete the documented node and cluster validation and recovery procedures after the VM environment is operational.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)