Rapid, Non-Disruptive Security Patching

VMware Cloud Foundation 9.1

August 2026

vmware by Broadcom

Table of Contents

Disclaimer....3

Executive Summary....4

The New Patching Reality....5

New Patching Capabilities in VCF 9.1....6

Knowing What to Patch....7

Protection Before a Patch Is Applied....8

- Virtual Patching via Distributed IDS/IPS for VMware vDefend 8
- Limiting Reachability with VMware vDefend Distributed Firewall 8

VCF's Layered Patching Architecture....9

- The Management Layer: VCF Operations and VCF Automation 9
- The Control Plane: vCenter, NSX, and Kubernetes 9
- The Data Plane: ESX, vSAN, and NSX Edge 10

Pre-Checks and Recoverability....13

De-Risking Workload Patching....14

Infrastructure Repaving as a Security Strategy....15

Conclusion....16

Additional Resources....17

Disclaimer

This document is intended to provide general guidance for organizations that are considering Broadcom solutions. The information contained in this document is for educational and informational purposes only. This document is not intended to provide advice and is provided "AS IS." Broadcom makes no claims, promises, or guarantees about the accuracy, completeness, or adequacy of the information contained herein. Organizations should engage appropriate legal, business, technical, and audit expertise within their specific organization for review of requirements and effectiveness of implementations.

Executive Summary

Exploitation can follow vulnerability disclosure quickly enough that a fixed quarterly patch schedule is insufficient for some vulnerabilities. VMware Cloud Foundation (VCF) 9.1 includes several patch and update methods that can reduce operator-visible interruption and workload impact. The result depends on the component, payload eligibility, redundancy, available capacity, and workload configuration.

VCF 9.1 uses different patching mechanisms for the management, control, and data plane layers. Organizations should assess each mechanism's effect on management services, data-plane forwarding, and virtual machine (VM) and container-based workloads. Optional VMware vDefend controls can reduce exposure while teams test a patch, but they do not replace remediation. Guest workload patching and infrastructure repaving require separate impact and recovery plans.

The New Patching Reality

Fixed quarterly patch cycles assume enough time for testing and a negotiated maintenance window. That assumption does not hold for every vulnerability.

Researchers, bug bounty programs, and automated scanners can shorten the interval between discovery and disclosure. Defenders still need time to assess exposure, test a fix, and obtain change approval.

The interval between publication and known exploitation can also be short. Zero Day Clock analyzes Common Vulnerabilities and Exposures (CVEs) with confirmed in-the-wild exploitation. For entries in its dataset, the reproducibility guide reports a median interval of 771 days for vulnerabilities published in 2018. The median was less than one day for vulnerabilities published in 2024. Vulnerabilities exploited on that timescale require a response process that can prioritize, test, and apply a fix outside a routine cycle.

The operational barriers remain familiar: change management, regression testing, maintenance windows, and limited capacity. A system without direct internet access can still be reachable through lateral movement, a compromised dependency, or an authorized account. Teams should assess actual reachability and compensating controls instead of treating "internal" as equivalent to isolated. They also need patch methods that reduce the cost and interruption of an urgent change.

VCF 9.1 coordinates patching across the management services, control plane, and data plane. Each layer uses mechanisms matched to its operational role, with different effects on management services, data-plane forwarding, and workloads. The degree of interruption depends on the patch method and environment.

New Patching Capabilities in VCF 9.1

VCF 9.1 includes capabilities designed to reduce patch duration or interruption. Their effect on running workloads depends on patch eligibility, component redundancy, available capacity, and workload configuration:

VMware vCenter Quick Patch: Quick Patch updates only the RPMs or binaries changed by an eligible payload. Quick Patch can reduce the total time for a vCenter patch operation by approximately 80%, from around 30 minutes to under five minutes, compared with traditional in-place patching. Management-service downtime depends on the affected services and can be zero.

VMware ESX Live Patching for Hosts with a Trusted Platform Module (TPM): For an eligible payload and compatible host, Live Patch applies a fix in memory without a host reboot or VM evacuation. If the payload updates the vmx runtime, compatible VMs briefly pause during Fast-Suspend-Resume (FSR).

Enhanced vCenter Reduced Downtime Upgrade (RDU): For vCenter 9.1.x and later payloads, RDU can obtain content from the online depot or use a mounted ISO in an offline environment. The final switchover typically requires less than five minutes of vCenter service downtime.

Rolling NSX Manager Updates: In a standard three-node cluster, Manager nodes are upgraded sequentially while at least two remain active. The management plane remains available, although users or API clients may need to reconnect. Existing data-plane forwarding continues during the Manager upgrade.

Coordinated ESX and NSX Host Updates: Updating the base ESX image also updates its embedded NSX host components. A Live Patch can include eligible NSX host components. Other payloads can require host evacuation and a restart, subject to cluster capacity and Edge redundancy.

Independent Patching for vSphere Supervisor and VMware vSphere Kubernetes Service (VKS): Platform teams can manage Supervisor and VKS cluster versions through rolling workflows. Application continuity depends on replicas, PodDisruptionBudget settings, and capacity for replacement nodes. A single-replica application can be interrupted while its pod is rescheduled.

VKS Fast Deploy: Fast Deploy can use Linked Mode or Direct Mode to provision replacement cluster nodes. It reduces provisioning time but does not change the application's availability requirements.

Knowing What to Patch

Rapid response begins before any patch is applied. VMware Security Advisories (VMSAs) identify affected products and versions, severity, vulnerability details, and available fixes. Security and infrastructure teams should monitor VMSAs and configure the notification options available through their Broadcom Support Portal accounts.

VCF Operations provides a centralized view of component versions, patches, and updates. In-product patch details identify affected services, eligibility, and expected impact when the payload supplies that information. Operators can use this inventory to map an advisory to deployed components, then select the supported remediation path for each component.

The VCF Operations Diagnostics Findings feature surfaces known conditions associated with security advisories, support cases, engineering findings, and customer-nominated issues. The Findings Catalog records details such as severity, associated knowledge base articles, and recommended actions. Operators can combine a finding's severity with reachability, business exposure, and remediation impact when setting patch priority. A finding supports that decision, but it does not replace review of the advisory and the affected configuration.

Protection Before a Patch Is Applied

A gap usually exists between disclosure and the completion of patch testing and rollout. VMware vDefend is an optional, separately licensed advanced service that can provide compensating controls during that interval.

Virtual Patching via Distributed IDS/IPS for VMware vDefend

Virtual patching uses an Intrusion Detection and Prevention System (IDS/IPS) signature or rule to detect or block traffic that matches a known exploit pattern. It does not modify the vulnerable software and should be treated as a temporary compensating control.

Teams must enable inspection and configure a policy with the intended detect or prevent action. Current content must also include an applicable signature. Prevent mode can block matching traffic without a guest reboot. Detect-only mode reports the match but does not block it. Teams must still test and apply the vendor-supplied software patch.

Distributed IDS/IPS for VMware vDefend places inspection in the distributed network path close to protected workloads. It inspects traffic selected by policy, not every flow by default. VCF 9.1 can also extend inspection to supported VKS pod traffic through its container networking integration. Policy scope, exemptions, traffic paths, and capacity determine which traffic is inspected.

For a known vulnerability, virtual-patch coverage depends on an applicable signature and current signature content. Connected, air-gapped, and egress-restricted sites need a documented process to obtain, stage, and validate content within their required response time. Custom signatures can address some gaps, but they also require validation and ongoing maintenance.

Behavioral detection evaluates activity such as anomalous DNS or beaconing patterns even when no signature fires. It can produce evidence for investigation. It is not a virtual patch for an unknown vulnerability and does not prove that every exploit attempt will be detected.

Use signature-based prevention and behavioral detection as separate, complementary controls. Their coverage depends on configuration, content, traffic visibility, and operating mode. Neither control substitutes for the permanent patch.

Limiting Reachability with VMware vDefend Distributed Firewall

VMware vDefend Distributed Firewall applies microsegmentation policy in the distributed network path for workloads in scope. Security teams can use that policy to restrict which systems may communicate with a vulnerable service while its patch is tested and staged. The control changes permitted network paths without changing the guest, but its effect depends on policy coverage, grouping accuracy, and allowed flows.

If one system is compromised, distributed firewall policy can restrict its permitted paths to other systems. Teams should validate the effective rules and required application flows before applying an emergency restriction.

These controls can reduce exposure between disclosure and remediation. They do not replace the permanent patch or eliminate the need to validate which traffic is protected.

VCF's Layered Patching Architecture

VCF manages infrastructure across three conceptual layers, each carrying distinct patching considerations and risk profiles.

- **The management layer** includes VCF Operations, VCF Automation, VCF Identity Broker, VCF Log Management, and VCF Operations for Networks.
- **The control plane layer** includes vCenter, NSX Manager, the vSphere Supervisor, and VKS clusters. These components govern compute, network, and container resources across the environment.
- **The data plane layer** consists of ESX hypervisors, vSAN storage, and NSX Edge and transport nodes: the infrastructure on which workloads actually run.

VMware Avi Load Balancer and VMware HCX use separate lifecycle workflows and require separate patch plans.

VCF 9.1 addresses each layer with mechanisms matched to its disruption profile. When a fix is eligible for a low-disruption method and the required redundancy and capacity are available, teams can apply it without taking workloads offline.

VCF 9 uses a unified product version and release model. The five release types are major, minor, maintenance, Express Patch, and Hot Patch. Express Patches are generally available, targeted releases for time-sensitive security or critical issues. Hot Patches are targeted fixes delivered to affected customers through Broadcom Support.

Release type helps an operator understand the scope of change, but it does not by itself establish the remediation method. Release notes and the in-product lifecycle view identify the affected component, dependencies, required sequence, and eligibility for mechanisms such as Quick Patch or Live Patch. Operators can then prioritize an applicable component patch without assuming that every fix requires a full-stack update.

The Management Layer: VCF Operations and VCF Automation

VCF 9.1 uses a declarative lifecycle model for VCF infrastructure components. Administrators select a target version in VCF Operations. Fleet Lifecycle plans the workflow, resolves dependencies, and obtains the required binaries. Operators remain responsible for reviewing the generated plan, prerequisites, impact, and recovery procedure before execution.

SDDC Lifecycle performs lifecycle operations for an individual VCF instance, and Software Depot stores installation, update, and patch artifacts. Air-gapped and partially connected sites must populate the applicable depot through the supported offline content process before lifecycle operations begin.

An update to VCF Operations or VCF Automation can interrupt the monitoring, lifecycle, portal, or API workflows that depend on the component being updated. Availability depends on the deployment model and update procedure. The Simple deployment uses single-node management appliances, while the High Availability deployment uses three-node clusters and is recommended for production. Before an update, operators should confirm the documented service impact and arrange any required out-of-band monitoring or operational access.

The management layer is architecturally separate from workload execution. Planned VCF Operations or VCF Automation service interruption during an update does not by itself stop running VMs or take vCenter or NSX services offline.

The Control Plane: vCenter, NSX, and Kubernetes

vCenter and NSX management form part of the control plane. An update can interrupt operations that depend on an affected service. VCF 9.1 provides rolling NSX Manager updates, vCenter Quick Patch with service-dependent downtime, and rolling workflows for Kubernetes cluster nodes. These mechanisms do not make every control-plane operation non-disruptive.

VMware vCenter Quick Patching

For a Quick Patch-compatible payload within a release line, vCenter Quick Patch updates only the RPMs or binaries modified by that payload. The affected services determine vCenter downtime, which can be less than one minute and, in some cases, zero. Traditional in-place patching updates every software package and can take the vCenter instance offline for an hour or more. During Quick Patch, VM and Kubernetes cluster deployments are not interrupted, and automation and API workflows continue.

Quick Patch availability is package-specific. The in-product patch details indicate whether a patch qualifies and identify the affected services, expected workload impact, and estimated downtime. If a patch is not compatible, use the update path specified for that release. The path can be a standard in-place update or Reduced Downtime Upgrade.

Reduced Downtime Upgrade for vCenter and Rolling Updates for NSX

For a supported version-level update, vCenter Reduced Downtime Upgrade (RDU) deploys a new appliance in parallel. It copies data and configuration while the source appliance remains operational. The final switchover stops the source services and starts the target services. The resulting interruption to vCenter services typically lasts less than five minutes; actual duration depends on the environment and update.

VCF 9.1 coordinates vCenter and NSX lifecycle steps, but the mechanisms remain distinct. RDU handles the vCenter appliance, while NSX Manager nodes use a rolling process. Operators should download the required content and run the available pre-checks before the maintenance window, then follow the dependency order in the generated update plan. The VCF 9.1 upgrade sequence is mandatory and varies with the starting environment.

NSX Manager nodes are upgraded sequentially. A standard three-node cluster keeps at least two nodes active, and transport-node assignments move to peer Manager nodes before a node is upgraded. The UI and API remain available during a supported rolling update, although users or API clients may need to reconnect as a node transitions. Existing data-plane forwarding continues while Manager nodes are updated. The management behavior of the new version becomes available after the NSX Finalize Upgrade step completes.

Lifecycle Management for VMware vSphere Supervisor and VMware vSphere Kubernetes Service Clusters

Containerized workloads add a separate lifecycle boundary. The vSphere Supervisor provides the Kubernetes control plane in VCF. VKS provisions and manages guest Kubernetes clusters on that infrastructure. Supervisor, cluster-node, and application updates each require their own plan.

The vSphere Supervisor coordinates rolling updates of VKS clusters. During a version upgrade, VKS replaces control-plane nodes first and then worker nodes. Health and compatibility pre-checks can stop the rollout if a node or cluster component is not ready. Changes to the cluster object are paused during the upgrade, but that management restriction does not by itself interrupt running applications.

VKS Fast Deploy uses Linked Mode or Direct Mode to provision replacement nodes during deployments and rolling upgrades. It can reduce rollout duration, but it does not change application-availability requirements.

Application availability during a rolling upgrade depends on replica placement, PodDisruptionBudget settings, and available capacity. An application can remain available when enough replicas stay ready and the disruption budget permits pod eviction. A single-replica application can be interrupted while its pod is rescheduled. The cluster needs capacity for replacement nodes and rescheduled pods. In an air-gapped VCF 9.1 environment, the required release content must be available through Fleet Depot Service before the upgrade begins.

The Data Plane: ESX, vSAN, and NSX Edge

Data-plane maintenance can affect workload execution, storage access, or network forwarding. Operators must evaluate each plane separately instead of inferring workload continuity from management-service availability.

ESX Live Patching

For a Live Patch-eligible payload and compatible host configuration, ESX Live Patch applies the fix in memory. The host does not enter maintenance mode, evacuate its VMs, or reboot. VMs remain powered on. If the payload updates the vmx runtime, compatible VMs undergo an FSR operation that briefly pauses guest execution. Live Patch supports TPM-enabled hosts. Depending on the payload, Live Patch can update the vmkernel, supported user-space daemons, NSX components, and vmx.

Operators may observe two behaviors during Live Patch, depending on the patch payload. If a user-space daemon such as hostd is patched and restarted, the host may briefly appear disconnected from vCenter. This expected transition does not affect running VMs. If the vmx runtime is patched, VMs undergo a Fast-Suspend-Resume (FSR) operation. FSR is a local, in-memory suspend and resume analogous to a vMotion within the same host. It keeps the VM powered on but briefly pauses

guest execution. vSphere also uses FSR for routine operations such as hot-adding virtual hardware. Traditional workloads typically do not experience a perceptible application pause during FSR. For vGPU-enabled VMs, FSR must also checkpoint GPU memory. VCF 9 reduces FSR stun time for these VMs. Internal testing with a VM configured with two 48 GB L40 GPUs measured an FSR time of approximately two seconds.

Some VM configurations are not compatible with the FSR step. VMs configured with Fault Tolerance, DirectPath I/O, or shared-disk clustering cannot undergo FSR, and neither can vSphere Pods. Their presence does not block the cluster operation. vSphere Lifecycle Manager reports the incompatible VMs and the reason before remediation. Those VMs continue on the pre-patch vmx runtime until they are power cycled or migrated to a patched host. Operators must track that follow-up remediation.

Live Patch applies to eligible Express Patches and Hot Patches; maintenance releases use standard remediation. Release notes and vSphere Lifecycle Manager identify eligibility before remediation. Do not infer eligibility from severity or release type alone.

Live Patch runs through vSphere Lifecycle Manager or its APIs, and pre-checks validate host resources. Hosts that use data processing units (DPUs) with vSphere Distributed Services Engine follow the standard remediation path.

Quick Boot

When a patch requires an ESX restart, Quick Boot can reduce host downtime on supported systems. It restarts the hypervisor without rebooting the physical hardware, bypassing Power-On Self Test (POST), firmware initialization, and device initialization. vSphere Lifecycle Manager checks Quick Boot compatibility before remediation and uses a regular reboot if the server or configuration is unsupported. Combined with VM migration managed by vSphere Distributed Resource Scheduler (DRS), Quick Boot shortens the maintenance window and the period of reduced cluster capacity.

Pre-Staging and Parallel Remediation

vSphere Lifecycle Manager can stage the desired ESX image to all or selected hosts before remediation. Staging transfers the image content to the target hosts without applying it. Moving the transfer ahead of remediation can reduce the active remediation phase and the period of reduced cluster capacity. It does not remove the time required for VM evacuation, image installation, a required restart, health checks, or return to service.

For clusters that support it, parallel remediation can shorten elapsed update time by patching multiple hosts concurrently. Administrators must place those hosts in maintenance mode and confirm that the remaining hosts can carry the workloads. vSphere Lifecycle Manager skips hosts not already in maintenance mode. Parallel remediation is not supported for vSAN clusters and cannot be combined with Live Patch. vSAN remediation and Live Patch process hosts sequentially.

DRS and vMotion Integration

When maintenance mode is required, DRS can use vMotion to evacuate running VMs to compatible hosts with sufficient capacity. VCF 9.1 can restrict evacuation to destination hosts that can satisfy the VMs' current compute demand. vMotion keeps each VM powered on but briefly stuns it during final switchover. Storage vMotion can move a running VM's virtual disks between compatible datastores and also includes a brief switchover stun. The effect depends on VM configuration, workload behavior, and available migration bandwidth. These capabilities allow many routine host and datastore maintenance operations to proceed without powering off the VMs.

vSAN Fault Domain-Aware Patching

In vSAN clusters, vSphere Lifecycle Manager remediates one host at a time. Sequential remediation alone does not establish storage-policy compliance. Objects can remain accessible while temporarily reporting reduced availability during maintenance. The result depends on the selected maintenance-mode data-migration option, object policy, capacity and fault domains, cluster health, and resynchronization state. Before remediation, run the vSAN data-migration pre-check, resolve inaccessible objects, and wait for any active resynchronization to complete.

NSX Edge and Transport Nodes

An NSX Edge cluster provides routing and other network services between virtual and physical networks. An Edge transition can affect traffic that depends on services provided by the cluster or network paths through the cluster.

NSX Edge nodes within a cluster are upgraded serially so peer nodes can continue forwarding traffic. Continuity still depends on Edge health, high-availability mode, routing convergence, and topology. A failed or delayed failover can interrupt an application path. Patching the ESX host that runs an Edge VM is a separate scenario, and its impact depends on Live Patch eligibility.

For a Live Patch-eligible payload, the ESX host remains out of maintenance mode and the Edge VM remains powered on. The remediation workflow therefore does not intentionally trigger an Edge shutdown or peer failover.

When Live Patch is unavailable, the host must enter maintenance mode. Edge Host Affinity and the Best Effort Restart (BER) policy cause pinned Edge VMs to shut down instead of migrating with vMotion. During vSphere Lifecycle Manager remediation, NSX checks peer Edge health before shutdown and blocks remediation if failover is not ready. A healthy peer assumes the affected forwarding functions according to the configured Edge high-availability mode and routing protocols. The interruption depends on the topology and protocol convergence settings. BER attempts to restart the Edge VM on a compatible host in its affinity group. If none is available, the VM remains powered off until its original host exits maintenance mode.

Pre-Checks and Recoverability

Pre-checks reduce avoidable failures, but they do not prove that an update will succeed. vSphere Lifecycle Manager checks compatibility and health before remediation. The VCF 9.1 RDU workflow runs a prerequisite check before the update, and Quick Patch identifies affected services and expected impact. Operators must resolve failed checks and still retain a component-specific recovery plan.

Recovery behavior differs by component. During RDU, the source vCenter remains available while data is copied. If the workflow fails, vCenter services can resume from the source appliance. That workflow recovery is not a substitute for a current vCenter file-based backup or a tested restore procedure. A failed VKS upgrade can stop until the unhealthy condition is corrected and the workflow is retried; VKS does not provide a general-purpose rollback guarantee. For ESX, vSphere Lifecycle Manager compares hosts with a desired image and remediates drift. The organization must validate the desired image before treating it as a recovery baseline.

De-Risking Workload Patching

Application owners remain responsible for patching guest operating systems, middleware, databases, and applications. Platform capabilities can support testing and recovery, but they do not make guest patching non-disruptive. Each application needs an impact assessment, a data-protection plan, and a rollback plan.

A pre-patch VM snapshot can provide a short-lived rollback option for some guest changes. It is not an independent backup or a guaranteed application recovery point because its delta disks depend on the base virtual disks. Quiescing can improve file-system or application consistency when the guest, VMware Tools, and application support it. If quiescing fails, the snapshot may be crash-consistent instead. Database recovery should use the database's native backup or a certified application-aware solution. Broadcom recommends retaining snapshots for no more than 72 hours, and its database snapshot guidance limits them to minutes or hours around a change. Before patching, test the rollback procedure and document the rollback point, potential data loss, consolidation capacity, and external dependencies.

A clone reproduces the VM configuration and disks at the time of cloning, but it does not reproduce every external dependency or production condition. Use the supported cloning workflow and isolate the clone before first power-on. Then change the host name and IP configuration as needed. Remove or update directory memberships, service identities, agents, credentials, and other environment-specific settings. A copied VM can retain a duplicate DHCP identity, and the default clone policy can copy virtual Trusted Platform Module secrets. VCF Automation can help provision repeatable test environments, but application owners must validate the test article and its isolation.

A guest patch and reboot can interrupt the affected VM and its application. DRS and vMotion support host maintenance; they do not provide application-level continuity for a guest reboot. Application replicas, failover, state synchronization, and health checks determine whether a service remains available. For protected traffic in scope, vDefend signature-based prevention and microsegmentation can provide temporary controls while application teams test the permanent patch.

Infrastructure Repaving as a Security Strategy

Incremental patching assumes that the existing infrastructure state is suitable for continued operation. After a confirmed compromise or loss of trust in foundational components, an organization may choose to rebuild infrastructure instead of patching it in place. Incident-response findings should determine the recovery scope.

Recommendation: Treat infrastructure repaving as an organization-defined recovery pattern. VCF Operations can automate deployment of workload domains. An organization can use that capability to build a replacement domain from approved inputs. Automation reduces manual deployment work, but it does not establish that the replacement environment is free of compromise.

Confidence in a replacement domain depends on trusted firmware, binaries, images, desired-state inputs, credentials, keys, network and storage services, and validation evidence. Workload movement also requires compatibility, spare capacity, dependency mapping, and an application-specific migration plan. Some migrations include a switchover pause, a restart, or a period of data synchronization. After a compromise, the incident-response team must decide whether moving workload state could also move malicious code, credentials, or persistence mechanisms.

Repaving can complement Live Patch, Quick Patch, rolling updates, backups, and incident response. An organization may also use planned rebuilds to reduce configuration drift. Repaving does not replace patching, credential recovery, forensic analysis, data recovery, or verification of external dependencies.

Conclusion

Some vulnerabilities are exploited too quickly for a fixed quarterly or monthly cycle to be the only response path. Organizations need a risk-based process for urgent patching when exposure and available evidence justify action outside the routine cycle.

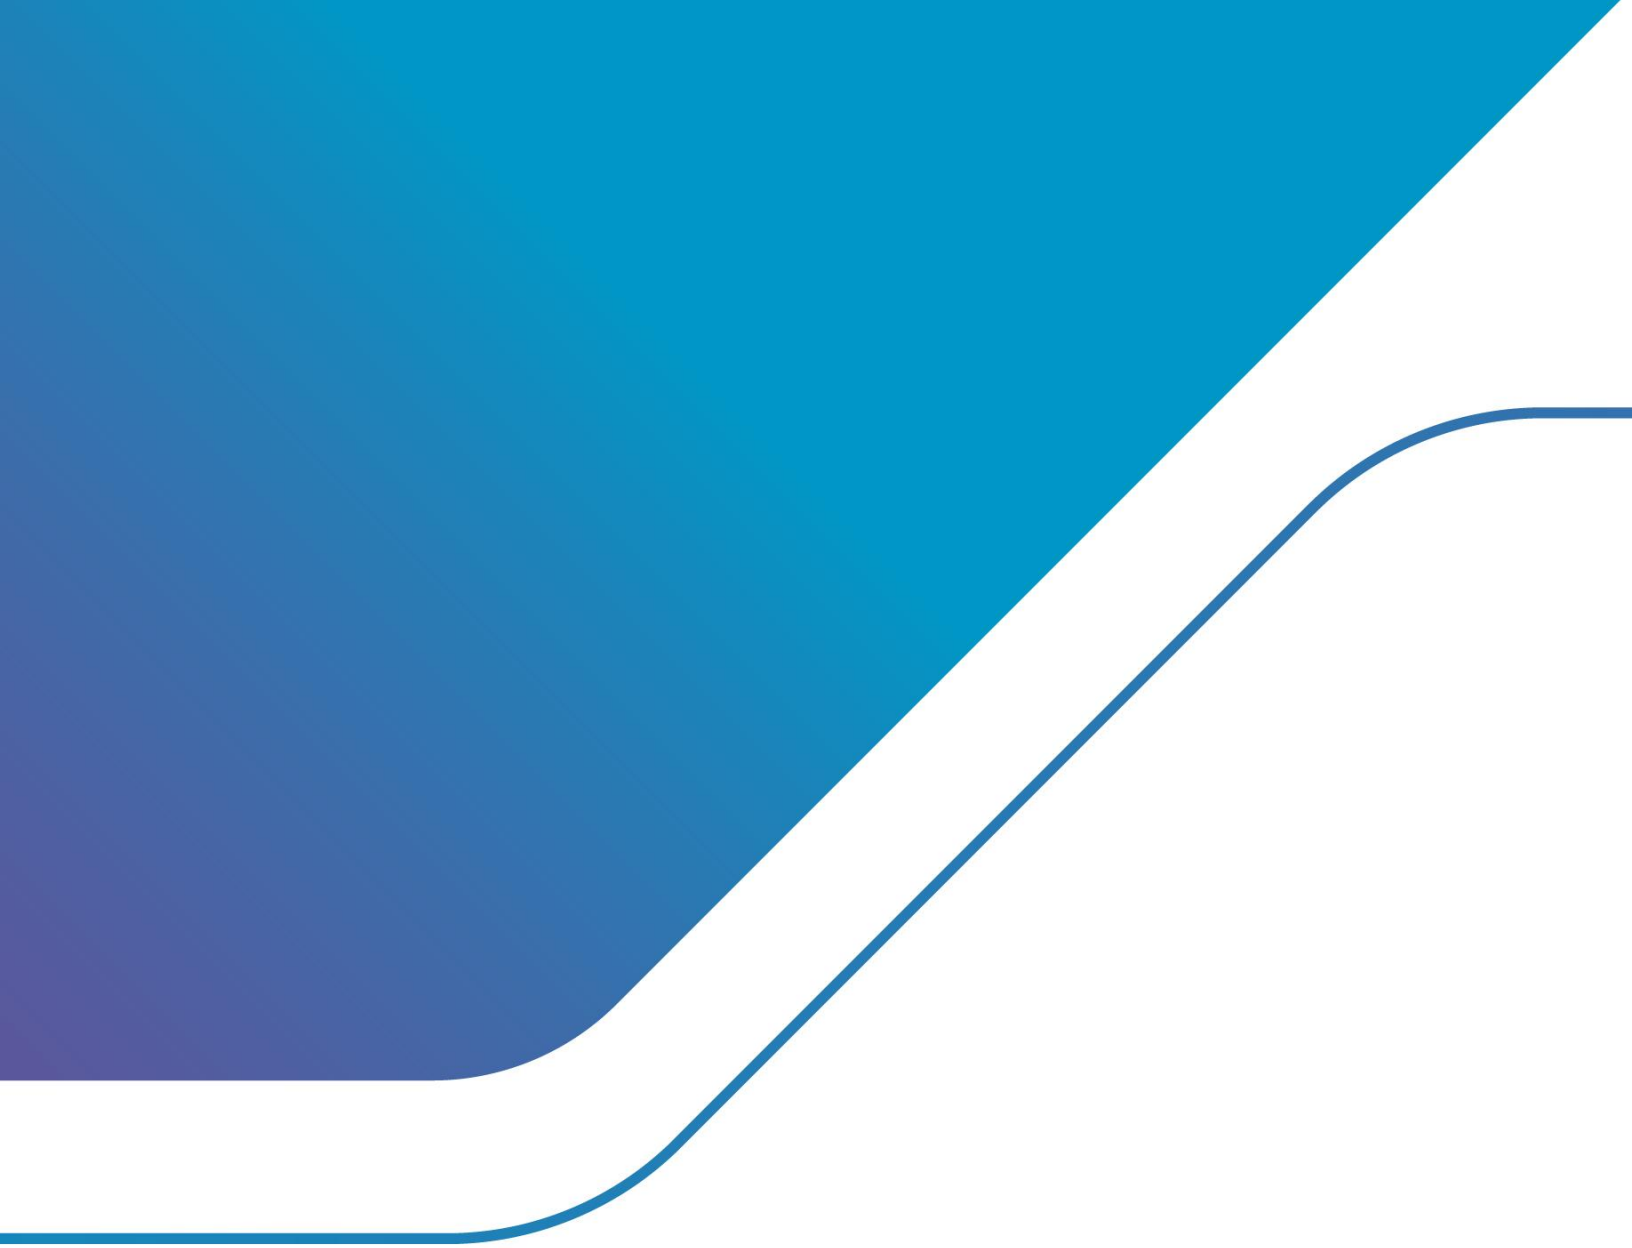
VCF 9.1 provides lower-disruption methods for eligible payloads. Quick Patch can reduce vCenter operation time, and vCenter service downtime depends on the affected services. RDU moves most vCenter update work before the switchover. Live Patch can avoid host evacuation and reboot, although an eligible vmx payload can briefly pause compatible VMs. Rolling NSX Manager and VKS workflows depend on redundancy, capacity, health, and application design.

Optional VMware vDefend controls can restrict permitted network paths, block traffic that matches applicable signatures, and report behavioral findings while a patch is tested. Their coverage depends on licensing, deployment, policy, content, traffic visibility, and operating mode.

For each component and workload, teams should define the eligible patch path, prerequisites, expected interruption by plane, validation checks, and recovery procedure. If a payload or configuration is not eligible for a lower-disruption method, use the supported standard remediation path.

Additional Resources

- Lifecycle Management Capabilities in VMware Cloud Foundation 9.1
  https://techdocs.broadcom.com/us/en/vmware-cis/vcf/vcf-service-administration-and-development/9-1.html
- VMware Security Advisories
  https://support.broadcom.com/web/ecx/security-advisory
- VMware Security Hardening Guidance
  https://github.com/vmware/vcf-security-and-compliance-guidelines/
- Identity and Access Management Best Practices for VMware Cloud Foundation 9.1
  https://brcm.tech/vcf-identity-design-best-practices
- VMware vDefend Design Library
  https://techdocs.broadcom.com/us/en/vmware-security-load-balancing/vdefend/design-library-for-vdefend/index.html



vmware® by Broadcom

Copyright © 2026 Broadcom. All rights reserved.
The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. For more information, go to www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.