Rapid Non-Disruptive Security Patching

VMware Cloud Foundation 9.1

June 2026

vmware by Broadcom

Table of Contents

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

The acceleration of vulnerability discovery, driven in part by AI-assisted security research, has fundamentally changed the patching calculus for enterprise infrastructure teams. Patch cycles that once operated on quarterly schedules must now accommodate responses within hours. VMware Cloud Foundation (VCF) 9.1 is engineered for exactly this reality: a layered patching architecture, orchestrated through the VCF Operations component, that enables operators to apply security fixes rapidly with minimal to no workload disruption across every tier of the stack.

This paper surveys the full spectrum of patching and update capabilities in VCF 9.1 that, together, keep VCF environments fully up to date while preserving the availability of virtual machine (VM) and container-based workloads.

It highlights the specific advances introduced in VCF 9.1 that make adopting this release a strategic security decision, examines how optional VMware vDefend security capabilities can provide interim protection in the critical window between vulnerability disclosure and patch availability, maps the platform's patching capabilities to the three distinct layers of the VCF stack, shows how the same de-risking discipline extends to patching the workloads themselves, and addresses infrastructure repaving as a complementary security strategy.

The New Patching Reality

Enterprise infrastructure security has entered a new operational era. The traditional model built on quarterly patch cycles, coordinated change windows, and staggered maintenance schedules was designed for a world where vulnerabilities were discovered gradually. That world no longer exists.

AI-assisted security research, expanded bug bounty programs, and the proliferation of automated scanning tools have compressed the window between vulnerability disclosure and active exploitation. Infrastructure teams that once had weeks to respond now need to respond in days and, in some cases, hours. Unpatched infrastructure is a compounding risk that grows with every hour of delay.

The same dynamic now applies on the adversarial side. In the wrong hands, AI can rapidly develop exploits that only months ago would have taken human attackers weeks to construct, if they could at all. Independent tracking bears this out: the Zero Day Clock project (zerodayclock.com), which measures time-to-exploit across thousands of CVEs with confirmed in-the-wild exploitation, reports that median time-to-exploit fell from 771 days in 2018 to approximately four hours in 2024. That leaves little margin between the moment a fix exists and the moment it is needed.

For many organizations, the structural barriers to rapid patching remain unchanged. Change management, regression testing, maintenance windows, and resource constraints do not dissolve because a Common Vulnerabilities and Exposures (CVE) flaw is published, and for systems perceived as internal or isolated, delaying a patch has often felt like the safer choice. That calculus is increasingly difficult to defend. Meaningful network isolation is rarely as complete as assumed: lateral movement from a perimeter breach, supply-chain compromises, and insider threats mean that "not internet-facing" no longer implies "low exploitation risk." The challenge is not whether to patch, but how to remove the operational cost that makes patching feel like the riskier option.

Broadcom, through its portfolio of VMware products, has built VCF 9.1 for this environment. VMware software delivers a coherent, integrated architecture for patching all layers of the software-defined data center, including the management services, control plane, and data plane, in a coordinated, disruption-minimizing manner. The premise of the VCF architecture is that speed and stability are not opposing forces: when the tooling is right, organizations can patch aggressively without introducing the downtime that once made rapid patching operationally unacceptable.

New Patching Capabilities in VCF 9.1

VCF 9.1 introduces advancements designed to minimize the operational friction and downtime traditionally associated with infrastructure patching. The following capabilities across the VCF stack enable a more responsive security posture with minimal impact on running workloads:

VMware vCenter Quick Patch: This new capability targets only the subsystems modified by a security or critical bug fix, reducing vCenter patch operation time by as much as 80%, in some cases with zero downtime for management services.

VMware ESX Live Patching for TPM-enabled Hosts: Starting in VCF 9.1, Live Patch coverage extends to hosts with Trusted Platform Modules (TPM), which are standard in current server generations, so eligible security fixes apply directly in memory with no host reboot and no VM evacuation.

Enhanced vCenter Reduced Downtime Upgrade (RDU): RDU now supports automated background downloads of update payloads and on-demand, non-disruptive pre-checks outside the maintenance window, shrinking the final switchover time to typically under five minutes.

Sub-second NSX Manager Updates: Orchestrated as rolling cluster operations that keep at least two nodes active throughout the process, NSX Manager upgrades in VCF 9.1 reduce network management downtime to sub-second durations.

Combined ESX and NSX Update Bundles: VCF 9.1 offers combined bundles that align the update sequence for both ESX and NSX components, which can reduce the total number of maintenance cycles and required DRS-led host evacuations.

vSphere-Independent Supervisor and VMware vSphere Kubernetes Service (VKS) patching: Platform teams can now upgrade vSphere Supervisor and VKS cluster versions independently of the broader VCF release cycle, applying Kubernetes-specific security patches through a rolling upgrade model that minimizes application disruption.

VKS Fast-Deploy via Linked Clones: VCF 9.1 uses linked clone technology to speed the deployment and upgrade of Kubernetes infrastructure, shortening the time to roll out security updates.

Knowing What to Patch

Rapid response begins before any patch is applied. Broadcom publishes VMware Security Advisories (VMSAs) that identify affected products and versions, describe the severity and nature of each vulnerability, and point to the releases that resolve it. Security and infrastructure teams can subscribe to advisory notifications so that evaluation starts the moment a vulnerability is disclosed, rather than when it surfaces in a periodic scan.

Within the environment, VCF Operations gives operators a fleet-wide view of component versions and applicable updates, so mapping an advisory to actual exposure does not require a manual, instance-by-instance inventory. Patch information presented in the product describes which services a given patch affects and what impact to expect, giving teams what they need to select the appropriate remediation path for each component and to prioritize by risk exposure rather than by operational convenience.

Beyond version visibility, VCF Operations Diagnostics Findings continuously evaluates the environment against known conditions, surfacing active findings across security vulnerabilities, CVE and VMSA exposure, issues with known resolutions, and trending support cases. Findings are drawn from security advisories, critical support cases, engineering-identified issues, and customer-nominated items, and a Findings Catalog provides a reference of all monitored conditions, including severity, associated knowledge base articles, and recommended actions. In practice, exposure assessment starts from targeted findings mapped to specific CVEs and affected components rather than from manual advisory tracking: a finding flagged as immediate against a deployed component, combined with a low patching cost for that component, is a clear and defensible basis for urgent action.

Protection Before a Patch Is Applied

Even with the most capable patching infrastructure, a gap always exists between the moment a vulnerability is disclosed and the moment a patch has been tested, staged, and applied across the entire environment. In a threat environment where exploits are developed within hours of CVE publication, that gap is a real and quantifiable attack risk. The optional VMware vDefend security capabilities available within VCF environments provide a critical layer of protection for precisely this window.

Virtual Patching via Distributed IDS/IPS for VMware vDefend

Virtual patching is a vulnerability-shielding tactic that creates a layer of security policy enforcement to intercept and block exploit attempts before they reach the vulnerable software. This mechanism, deployed through an Intrusion Detection and Prevention System (IDS/IPS), is a temporary compensating control that helps prevent the exploitation of known vulnerabilities without modifying the underlying application code.

The primary benefit of virtual patching is that it gives security teams faster protection against threats and can usually be deployed anytime with no maintenance window or workload disruption required. This provides the time needed to test, stage, and apply a permanent vendor-supplied patch while maintaining continuous service availability.

Distributed IDS/IPS for VMware vDefend performs virtual patching by embedding its distributed IDS/IPS engine directly into the hypervisor fabric. This placement allows the IDS/IPS to inspect all network traffic at the virtual NIC (vNIC) of every workload, forming a proximity shield at the last network hop before the workload itself. Through automated dynamic policies and frequently updated signatures, vDefend can provide immediate protection against both external and lateral threats.

Alongside signatures, the Distributed IDS/IPS includes behavioral analysis capabilities for detecting unusual traffic patterns that may indicate targeted or zero-day exploits for which no signature yet exists.

For a vulnerability that has been publicly disclosed but not yet patched, this combination of signature and behavioral detection provides a meaningful detection and prevention capability that can alert operators to active exploitation attempts before a patch is in place.

Distributed IDS/IPS for VMware vDefend forms a defense-in-depth layer that operates continuously and independently of patch status. It is not a substitute for patching (no security control is), but it materially reduces the risks associated with the gap between vulnerability disclosure and permanent patch deployment.

Limiting Reachability with VMware vDefend Distributed Firewall

A disclosed vulnerability is only exploitable by an attacker who can reach it. VMware vDefend Distributed Firewall enforces microsegmentation policy at the vNIC of every workload, allowing security teams to restrict which systems can communicate with a vulnerable service while its patch is tested and staged. Tightening policy around affected workloads requires no change to the workloads themselves and no maintenance window, and it can reduce the pathways available to exploit a vulnerability from the entire environment to a narrow, deliberately chosen set.

Microsegmentation also addresses the scenario that makes pre-patch windows most dangerous: lateral movement. An attacker who compromises one system during the exposure window must still traverse the network to reach higher-value targets, and distributed firewall policy restricts that traversal regardless of the patch status of the systems on either end.

For organizations adopting VCF 9.1 in a high-velocity threat environment, enabling these capabilities alongside VCF's patching architecture provides a layered security posture that addresses both the known (patched vulnerabilities) and the emerging (pre-patch exposure windows) risks.

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 Operations for Logs, and VCF Operations for Networks.
- **The control plane layer** encompasses vCenter, NSX Manager, the vSphere Supervisor, and the VMware vSphere Kubernetes Service (VKS) clusters: the components that govern compute, network, and container resource management 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.

This paper concentrates on these core layers of the VCF stack. Components with their own lifecycle workflows, such as VMware Avi Load Balancer and VMware HCX, follow update processes documented separately and are outside the scope of this discussion.

VCF 9.1 addresses each layer with purpose-built mechanisms. The result is not a single monolithic patching workflow but a coordinated set of capabilities that match the disruption profile of each tier. An urgent CVE no longer always forces a choice between security response and workload availability.

These capabilities rest on a unified release model shared across VCF and vSphere, which replaces the prior ESX-specific versioning scheme and user-facing aliases such as "Update 2a" or "U3b." The five release types are major, minor, maintenance, Express Patch, and Hot Patch. For major and minor releases, ESX and NSX must be upgraded together; for maintenance releases, Express Patches, and Hot Patches, ESX and NSX can be updated independently. The distinction matters operationally because the release type determines a fix's eligibility for low-disruption mechanisms such as ESX Live Patch.

Patch releases within each VCF 9 layer are also designed for surgical precision. They address time-sensitive security and critical issues between major, minor, and maintenance releases, and do not require a synchronized update across all components or a new Bill of Materials implementation. Administrators can select and apply patches to individual components based on risk exposure and operational cost, giving teams the flexibility to prioritize the most critical fixes without triggering a full-stack update cycle.

The Management Layer: VCF Operations and VCF Automation

Starting in VCF 9.1, the VCF infrastructure components use a declarative lifecycle management model. Rather than executing discrete patching commands, administrators define a target version for the VCF environment. The Fleet Lifecycle service in VCF Operations orchestrates the entire upgrade or patch process across the fleet, including planning, dependency resolution, and binary procurement, making the patch process faster and more predictable. This approach reduces the manual steps that have historically made management-layer patching labor-intensive and error-prone.

Complementing the Fleet Lifecycle service, the SDDC Lifecycle service handles installation, update, and patch operations at the individual VCF instance level. The Software Depot, also new in VCF 9.1, is an internal binary store for installation, upgrade, and patch artifacts, and can operate at either the fleet level or the instance level, so air-gapped and partially connected sites follow the same declarative patching workflow as connected ones rather than a separate manual process.

While VCF Operations updates itself, monitoring, alerting, and fleet management capabilities are temporarily unavailable, regardless of deployment configuration. For VCF Automation, the update is orchestrated across the underlying Kubernetes cluster on which its services run: uninterrupted service through the update requires the three-node High Availability deployment model recommended for production environments. In the single-node Simple deployment model, intended for proof-of-concept, development, and testing scenarios, VCF Automation services are unavailable during the update window.

The management layer is architecturally separate from the workload plane, so these update windows carry no workload risk: an issue encountered during a VCF Operations or VCF Automation update has no impact on running VMs, vCenter, or NSX.

The Control Plane: vCenter, NSX, and Kubernetes

vCenter and NSX management form the control plane, and any downtime during their update can ripple through every operation that depends on them. VCF 9.1 holds that disruption to a minimum through several complementary mechanisms: sub-second VMware NSX Manager updates, near-zero-downtime vCenter Quick Patching, and disruption-minimizing rolling upgrades for Kubernetes workloads.

VMware vCenter Quick Patching

vCenter Quick Patch is the recommended method for applying security and minor maintenance patches within a release line. Rather than updating every software package on the appliance regardless of whether it changed, Quick Patch targets only the specific RPMs or binaries modified in the patch payload. This narrows the maintenance window sharply: depending on which vCenter services the patch affects, downtime is brief and in some cases zero. Throughout the operation, VM deployments, Kubernetes cluster operations, automation, and API workflows continue without interruption.

Not every patch is Quick Patch compatible; eligibility depends on the patch payload. The patch details that are exposed in-product indicate whether a given patch qualifies and disclose the affected services, expected workload impact, and estimated downtime before the operator commits to the operation. For patches that fall outside Quick Patch scope, such as OS-level changes or version transitions, Reduced Downtime Upgrade is the appropriate path. By contrast, traditional in-place patching updates every appliance package regardless of change and can take a vCenter instance offline for an hour or more, which is the cost Quick Patch is designed to remove.

Reduced Downtime Upgrade for vCenter and NSX

For version-level upgrades that cannot be addressed through Quick Patching, vCenter Reduced Downtime Upgrade (RDU) uses a migration-based approach: a new vCenter appliance is deployed in parallel, and all data and configurations are copied from the source appliance while it remains fully operational. The preparation phase is the most time-consuming portion of the process; however, it requires no downtime. The only service interruption occurs during the final switchover when the source appliance is stopped and the new instance takes over, which takes approximately five minutes under normal operating conditions.

In VCF 9.1, the RDU workflow is formally integrated with NSX management update orchestration. The preparation stage can be completed entirely before entering the maintenance window; within the window, NSX management nodes are updated first, followed by the vCenter instance switchover. This sequencing both minimizes the maintenance window footprint and updates the network control plane and compute management plane in a coordinated, dependency-aware order.

During the NSX management node updates themselves, the network management plane remains continuously available. Manager nodes are upgraded sequentially, with each node's transport node assignments redistributed to peer nodes before it is upgraded, so at least two manager nodes stay active throughout, and API calls, configuration changes, and transport node operations continue uninterrupted. UI and API behavior remains consistent with the pre-upgrade version until the Finalize Upgrade step, exposed in both the NSX and VCF Operations interfaces, is complete, at which point the new version's capabilities become available.

VMware vSphere Supervisor and VMware vSphere Kubernetes Service Cluster Lifecycle

VCF environments running containerized workloads introduce an additional dimension of lifecycle management. The vSphere Supervisor is the control plane for Kubernetes workloads within VCF, and the VMware vSphere Kubernetes Service (VKS) provisions and manages the guest Kubernetes clusters running on that infrastructure. Both require their own lifecycle management discipline, and both are addressed within the VCF patching architecture.

The vSphere Supervisor and VKS clusters use a rolling update model that minimizes downtime for cluster workloads during the update process. Rolling updates encompass Kubernetes software version upgrades, as well as the infrastructure and services supporting those clusters, including virtual machine configurations, services, namespaces, and custom resources. The system enforces compatibility pre-check conditions before updates proceed and supports rollback if a cluster upgrade encounters a failure, preserving workload continuity throughout the process.

VKS version upgrades are managed at the vSphere Supervisor level and propagate through clusters in accordance with the rolling update model, maintaining availability for containerized applications throughout the process. VCF 9.1 also shortens the rollout itself: VKS Fast-Deploy uses linked clone technology to provision the new cluster nodes created during deployments and rolling upgrades, reducing the time required to move a security update across a large population of clusters.

Workload availability through a rolling upgrade depends in part on application configuration. Applications running multiple replicas typically experience no interruption, while single-replica deployments may see brief disruption as pods are evicted and rescheduled onto nodes running the new version. The cluster must also have sufficient capacity to absorb rescheduled workloads while nodes are being replaced. Disconnected environments should plan for additional steps as well: air-gapped sites must stage Kubernetes release artifacts locally before remediation can proceed.

The Data Plane: ESX, vSAN, and NSX Edge

The data plane is where workloads actually run, so it carries the strictest disruption constraints of any tier. The VCF architecture is built to minimize impact on running virtual machines and containers during data plane maintenance.

ESX Live Patching

The most significant advancement for data plane patching is the expansion of ESX Live Patch. Live Patch applies security and urgent bug fixes directly to running memory, leaving virtual machines on the host fully operational, with no maintenance window, no VM evacuation, and no host reboot. Starting with VCF 9.1, Live Patch is supported on TPM-enabled hosts, extending in-memory patching to current generations of server hardware. The surface area eligible for Live Patch is broad, covering the vmkernel itself, user-space daemons, NSX components, and the virtual machine execution runtime (vmx).

Depending on the patch payload, a Live Patch operation can surface two operator-visible behaviors. If a user-space daemon such as hostd is patched and restarted, the host may briefly appear disconnected from vCenter; this is expected and does not affect running VMs. If the vmx runtime is patched, VMs undergo a Fast-Suspend-Resume (FSR) operation: a local, in-memory suspend and resume analogous to a vMotion within the same host. FSR is non-disruptive and is already used in routine vSphere operations such as hot-adding virtual hardware. FSR performance for vGPU-enabled VMs has been significantly improved in VCF 9, allowing clusters hosting large AI/ML workloads to be Live Patched without application disruption.

Some VM configurations are not compatible with the FSR step of a Live Patch operation. VMs configured with Fault Tolerance or DirectPath I/O, vSphere Pods, and VMs participating in shared-disk clustering configurations cannot undergo FSR. Their presence on a host does not block Live Patch operations: vSphere Lifecycle Manager compliance scans identify and report incompatible VMs, along with the reason for the incompatibility, before remediation starts on the cluster.

Live Patch applies to Express Patches and Hot Patches; maintenance releases require standard remediation. Patch release notes and the vSphere Lifecycle Manager interface identify Live Patch eligibility before remediation begins, so operators know in advance whether a given patch can be applied in memory or requires a maintenance-mode workflow.

Quick Boot

When a patch does require a host reboot, Quick Boot, a mature ESX capability available since vSphere 6.7, sharply reduces the time the host spends offline. Quick Boot skips hardware initialization during the restart cycle and reloads only the hypervisor, compressing what would otherwise be a multi-minute hardware Power-On Self Test (POST) sequence into a fraction of that time. Combined with DRS-driven VM migration, Quick Boot reduces both the maintenance window duration and the period of reduced cluster capacity.

Pre-Staging and Parallel Remediation

vSphere Lifecycle Manager supports pre-staging of ESX images: the process of downloading and preparing all software and firmware components to the target hosts before the maintenance window begins. Staging is performed in parallel across all hosts in a cluster by default, eliminating the download phase from the live maintenance window and reducing the time each host spends offline to the minimum necessary for the actual update.

For environments where cluster resources permit multiple simultaneous host operations, parallel remediation allows multiple hosts within a cluster to be patched concurrently based on available capacity. The combination of pre-staging and parallel remediation can reduce total patching duration for a large cluster from hours to a fraction of that time.

DRS and vMotion Integration

When maintenance mode is required, vSphere's Distributed Resource Scheduler (DRS) and vMotion work in concert to evacuate workloads from hosts before remediation begins. vMotion migrates running virtual machines live, with no disruption to users or loss of service, to other hosts in the cluster. From the application's perspective, the underlying host maintenance is invisible. Storage vMotion provides the equivalent capability for storage-layer maintenance, so neither compute nor storage operations require application downtime for routine patching.

vSAN Fault Domain-Aware Patching

In vSAN clusters, vSphere Lifecycle Manager applies fault domain awareness during remediation: hosts within a single fault domain are patched sequentially, and remediation progresses one fault domain at a time. This sequencing maintains vSAN data availability throughout the patching process.

NSX Edge and Transport Nodes

NSX Edge nodes are the boundary between physical and virtual networks: any interruption affects traffic forwarding for all workloads depending on that connectivity, making disruption minimization essential during patching.

Two distinct maintenance scenarios apply. Updating the NSX Edge software itself is a rolling operation that VCF orchestrates across the members of an Edge cluster: each node is updated and returned to service before the next begins, and traffic forwarding fails over to cluster peers for the duration of each node's update. The other scenario is patching the ESX host on which Edge node VMs run, where the impact depends on whether the patch supports Live Patch.

For patches that support Live Patch, the host is patched in memory without entering maintenance mode, requiring no Edge node shutdown and no protocol-level failover.

When an update cannot support Live Patch and a host running Edge node VMs must enter maintenance mode, Best Effort Restart policies and Edge host-group affinity maintain continuity without vMotioning Edge node VMs. Protocol-level failover between Edge cluster peers transfers traffic forwarding responsibility in far less time than a live migration would require, minimizing the network interruption window.

Host-group affinity enforces placement boundaries for Edge node VMs and validates peer readiness before vLCM initiates maintenance mode and the Edge VM is shut down, so Edge failover proceeds cleanly.

Pre-Checks and Recoverability

Change-related risk, not doubt about the fix itself, is what has historically led organizations to defer patching. The VCF patching architecture addresses that risk directly. Before a change is committed, the platform verifies that it is likely to succeed: vSphere Lifecycle Manager runs compatibility and health pre-checks before remediation begins, RDU offers on-demand, non-disruptive pre-checks that run outside the maintenance window, and Quick Patch discloses affected services and expected impact before the operator commits to the operation.

When a change does go wrong, the architecture favors recoverability. RDU's migration-based design keeps the source vCenter appliance intact until the final switchover, preserving a fallback position through the riskiest portion of the upgrade. vSphere Supervisor and VKS rolling updates support rollback if a cluster upgrade encounters a failure. And because ESX host state is defined declaratively, the desired-state image managed by vSphere Lifecycle Manager is itself a recovery tool: a host that drifts or fails validation can be remediated back to a known-good image rather than debugged in place.

De-Risking Workload Patching

The compression of disclosure-to-exploit timelines does not stop at the infrastructure. The same pressure applies to the software running inside virtual machines: guest operating systems, middleware, databases, and applications. Patching that software is the application owner's responsibility rather than the platform's, but the platform has a direct role in removing the operational risk that makes application teams defer those patches, and VCF provides workload-level equivalents of the same safeguards: verification before a change is committed, and recoverability when one goes wrong.

The first is the restore point. A VM snapshot taken before a guest patch captures the disk, virtual machine configuration, and, optionally, memory state of the running system, so a patch that fails or destabilizes the application can be reverted in minutes rather than recovered through rebuild or restore. Snapshots are a short-lived safety net rather than a backup strategy, but for the duration of a patch window they convert an irreversible change into a reversible one, which is precisely the property that makes rapid patching defensible to a change advisory board.

The second is the faithful test article. Cloning a production VM produces an exact copy of the system as it actually exists, accumulated configuration and all, and attaching that clone to an isolated NSX network segment allows a patch to be rehearsed against the real system without touching production or risking duplicate-identity conflicts. Where teams need this repeatedly, VCF Automation can provision standing test environments on demand, making patch rehearsal a routine step rather than a special project.

The remaining pieces are already in place. DRS and vMotion keep neighboring workloads unaffected while a guest is patched and rebooted, and containerized applications inherit the VKS rolling update model. The pre-patch protections also extend here: the vDefend virtual patching and microsegmentation capabilities shield vulnerable software regardless of whether it is a platform component or an application inside a guest, buying application teams the same test-and-stage time the platform's own patching enjoys.

Infrastructure Repaving as a Security Strategy

Incremental patching, even when executed efficiently at every layer, operates on the assumption that the existing infrastructure state is a sound foundation for continued operation. In some security scenarios, particularly after a confirmed compromise, a prolonged exposure window, or a critical vulnerability affecting foundational components, an organization may conclude that the most reliable path to a known-good state is not patching in place but redeploying from scratch.

VCF's workload domain architecture is designed to support exactly this pattern, sometimes referred to as infrastructure repaving. A workload domain in VCF is a logical, self-contained pool of compute, network, and storage infrastructure, provisioned and managed through VCF Operations. Because workload domains are defined declaratively and deployed through automated processes, a new domain can be provisioned, validated, and placed into service at infrastructure scale without manual configuration steps.

This means that redeploying a workload domain to a pristine, fully patched state is not a contingency procedure that requires heroic manual effort; it is an automated operation within the normal operational envelope of VCF. ESX hosts are reimaged to the current desired-state image managed by vSphere Lifecycle Manager, vCenter and NSX are redeployed to current versions, and the domain is returned to a verifiable clean state. Workloads can be migrated to a new domain before the original is torn down, preserving availability throughout the transition.

Infrastructure repaving complements in-place mechanisms such as Live Patch, Quick Patch, and rolling upgrades. For environments concerned about persistent threats, supply-chain vulnerabilities affecting deep system components, or the integrity of long-running infrastructure that predates improved patching practices, periodic repaving provides a level of assurance that patch-on-patch operations cannot: confidence that the infrastructure is running exactly and only the software it was deployed to run.

Conclusion

The rate of vulnerability discovery is not returning to the pace of prior years. Organizations that continue to operate on the assumption that quarterly or monthly patching cycles are adequate are compounding risk with each cycle. The infrastructure platform must be capable of keeping pace with the threat environment, not lagging behind it.

VMware Cloud Foundation 9.1 delivers that capability. Through a layered architecture covering every tier of the software-defined data center, from declarative lifecycle management of the management plane, to near-zero-downtime vCenter Quick Patching and RDU-based control plane upgrades, to zero-maintenance-window ESX patching at the data plane, VCF 9.1 enables security-responsive operations without compromising workload availability. Infrastructure repaving through workload domain redeployment provides an additional strategic option for achieving verifiable clean-state infrastructure at scale.

The optional VMware vDefend security capabilities extend protection into the pre-patch exposure window, providing detection and lateral movement control while patches are staged and applied.

For infrastructure organizations working to match patching velocity to today's threat environment, VCF 9.1 removes the operational cost that has long forced a trade-off between security response and workload availability, so teams can patch at the speed the threat environment now demands.

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/
- VMware vDefend Design Library
  https://techdocs.broadcom.com/us/en/vmware-security-load-balancing/vdefend/design-library-for-vdefend/index.html