Delivering a Unified Cloud Experience with VMware Cloud Foundation

Key components of a unified cloud experience

- Self-service and on-demand infrastructure provisioning
- End-to-end automation and orchestration
- Seamless integration of DevOps practices
- Support for modern application development
- Unified user interface and APIs
- Scalability and flexibility
- Built-in security and compliance

With a unified cloud experience, your application teams can consume all necessary tools, infrastructure resources, and services through a single, integrated platform. This enables them to build, test, deploy and manage all kinds of applications—including traditional apps, modern apps, cloud-native apps, and AI/ML apps—without having to deal with the complexities of infrastructure management. However, delivering a unified cloud experience poses several challenges that can affect efficiency and productivity:

- **Lack of automation and orchestration capabilities**
- **Infrastructure silos and interoperability challenges**
- **Inadequate support and skills for modern app development**
- **Limited infrastructure visibility and monitoring**
- **Security and privacy concerns due to outdated infrastructure**

The solution: VMware Cloud Foundation

VMware Cloud Foundation (VCF) is a comprehensive private-cloud platform that combines the scale and agility of public cloud with the security and performance of private cloud to offer industry-leading TCO. Delivering integrated, enterprise-class compute, networking, storage, management and security across all endpoints, VCF abstracts away the complexity of underlying infrastructure, giving your platform teams self-service access to infrastructure resources for on-demand provisioning and management without IT intervention.

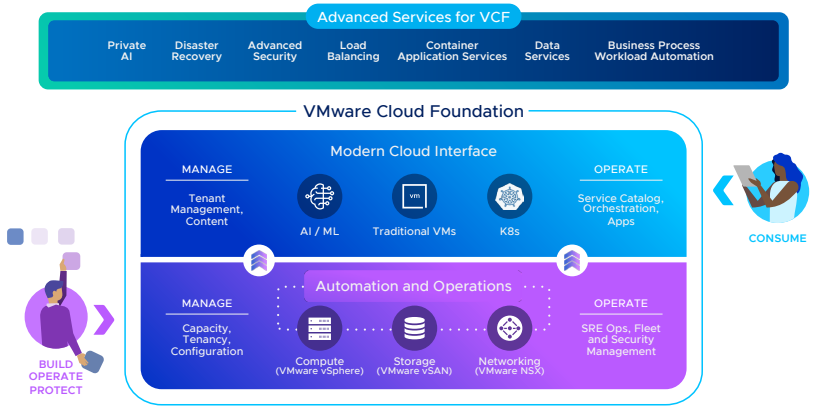

Figure 1: Deliver a unified cloud experience with VMware Cloud Foundation.

vmware® by Broadcom

Key benefits of using VCF for running modern workloads

- Maximized value and cost savings
- Faster time to market for modern apps
- Improved developer productivity
- Better operational efficiency
- Enhanced infrastructure scalability, consistency and reliability

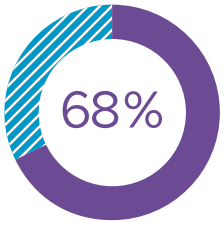

faster to bring new applications to the market

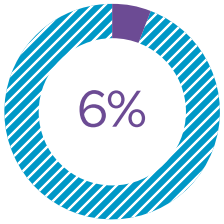

improved developer productivity

Source: IDC White Paper, sponsored by VMware by Broadcom, The Business Value of VMware Cloud Foundation, doc #US52312224, August 2024)

"We are transitioning all our software to modern, containerized apps. With VMware, we have progressed quickly on that journey, and transforming our software has helped us become a more agile business moving in the direction consumers want."

Vishwas Chitale
CEO & CTO, Chitale Dairy

Provide a centralized consumption experience

VCF Automation, an integral component of VCF, introduces an API abstraction layer that seamlessly integrates with the various components of your infrastructure. It creates a centralized consumption experience, enabling users to manage resources from multiple VCF instances as if they were interacting with a single cohesive platform. This means your application teams can now interact with one central API front, simplifying the provisioning of IaaS.

VCF Automation delivers centralized IaaS consumption through Infrastructure as Code (IaC), a methodology that empowers teams to curate blueprints. Utilizing YAML and Kubernetes manifests alongside various scripting languages, IaC becomes the vital conduit through which your organization can provision and manage both modern and traditional workloads.

Use case: Build, run and manage Kubernetes and other modern workloads

VCF provides a modern, flexible infrastructure to build, run and manage modern workloads such as containers, micro-services, big data analytics and high-performance computing workloads. Here are some key features and capabilities:

- **Natively integrated Kubernetes runtime** – With the inclusion of VMware vSphere Kubernetes Service (VKS), VCF delivers native, CNCF (Cloud Native Computing Foundation)-compliant Kubernetes container orchestration directly into the platform, at no extra cost.
- **Faster delivery of the latest Kubernetes versions** – With VKS as an independent service, VCF decouples Kubernetes releases from vCenter, enabling asynchronous updates aligned with upstream Kubernetes.
- **Ability to run Kubernetes workloads natively on ESX hosts with VMware vSphere Supervisor** – By creating a Kubernetes control plane directly on the hypervisor layer, vSphere Supervisor enables provisioning of IaaS services such as Network Service, Storage Service, VM Service and vSphere Kubernetes Service. It allows administrators to create and configure vSphere Namespaces with allocated resources for DevOps teams. DevOps engineers can then run Kubernetes workloads, deploy Kubernetes clusters with VKS, and manage containers or VMs on the same platform, while administrators maintain full visibility and control over resource usage and deployments via VCF Automation.
- **Easy consumption of IaaS services** – VCF Automation Services (previously known as Supervisor Services) are vSphere-certified Kubernetes operators that deliver IaaS components and tightly-integrated Independent Software Vendor services to developers. By installing VCF Automation Services on vSphere Supervisor, administrators can extend the platform with additional IaaS services such as VKS, VM Service, Network Service, Storage Services, Capacity Planning, and other supervisor services such as Object Storage, Image Registry (Harbor), and backup and recovery service.
- **Kubernetes automation** – VCF Automation allows IT admins to automate the configuration and deployment of Kubernetes clusters and Supervisor Namespaces using IaC templates while maintaining governance and control.

Key benefits of using VCF to build, run and manage Private AI

- Enhanced data privacy and security
- Accelerated performance of GenAI models
- Simplified GenAI deployment
- Improved cost efficiency
- Greater control and customization

"VCF and Private AI Foundation allowed us to quickly and efficiently implement AI technologies in our applications in a secure manner, whether it's for knowledge-based chatbots assisting staff or delivering financial answers directly to our members across multiple devices. It helped us automate everyday transactions, streamline loan processing, and proactively identify member needs."

Mark Fournier
Chief Information Officer,
US Senate Federal Credit Union

- **Resource policies with Policy as Code** – With CNCF-compliant VKS built into the VCF platform, IT admins can build custom policies for VMware Kubernetes clusters and other IaaS resources, using new YAML-based Policy as Code without having to rely on external tools.
- **End-to-end visibility of Kubernetes clusters** – VCF Operations Management Pack for Kubernetes enables users to visualize, monitor and troubleshoot Kubernetes clusters and their components, providing end-to-end visibility.
- **Comprehensive data services management with VMware Data Services Manager** – DSM enhances data service delivery and management for AI/ML, cloud-native, and traditional applications, and enables IT to offer self-service, enterprise-hardened Postgres, MySQL, MinIO AIStor and Microsoft SQL Server to their teams, boosting developer innovation, reducing IT costs, and ensuring data resilience.
- **Consistent networking for VMs and containerized workloads** – VMware NSX delivers advanced container networking with granular network policies applied at the per-container level. It simplifies deployment with native container-to-container networking and provides end-to-end visibility and troubleshooting for all workloads.
- **Container security with Antrea CNI** – VMware NSX offers full enterprise support for Project Antrea, a CNCF sandbox project. It delivers a simplified network interface that allows you to implement Kubernetes networking using Open vSwitch (OVS), enabling easy pod-to-pod communication, enforcing network policies, encrypting traffic, and enhancing observability.
- **Advanced management of containerized workloads** – VMware Tanzu Platform, an advanced service that can be purchased separately, provides advanced capabilities such as container management, enterprise app networking, local build environments, application middleware, secrets management, and hardened OSS images, eliminating costly investments in DIY app platform solutions.

Use case: Build, run and manage Private AI

VCF enables you to build, run and manage private AI models and applications in your environment while maintaining control over sensitive data and meeting stringent security, compliance and privacy requirements. Here are some key features and capabilities:

- **Simple and secure GenAI deployment with VMware Private AI Foundation with NVIDIA** – This joint GenAI platform enables your enterprise to run RAG workflows, fine-tune and customize LLM models, and run inference workloads in your data centers, addressing privacy, choice, cost, performance and compliance concerns. An advanced service that needs to be purchased separately on top of VCF, it offers the following key capabilities:
  - **Deep learning VM templates** – These templates come pre-configured with required software frameworks like NVIDIA NGC, libraries, and drivers, saving users from the burden of setting up each component.
  - **Vector databases for enabling RAG workflows** – For fast querying of data and real-time updates to enhance the outputs of the LLMs, VMware has enabled vector databases by leveraging pgvector on PostgreSQL.

"VMware Cloud Foundation provides everything we need: compute, storage and networking as well as load balancing and flexibility. With VMware Aria Automation we're providing a public cloud-like user experience while maintaining security, compliance and control. We are a small team, but VMware enables us to offer anything as a service (XaaS)."

Philippe Morel
Director of IT Operations and Infrastructure, EPFL

- **Catalog setup wizard** – This capability enables admins to efficiently design, curate and provide optimized AI infrastructure items in a self-service catalog through the VCF Modern Cloud Interface. After publication, DevOps and data scientists can effortlessly access the GPU-enabled deep learning VMs and Kubernetes clusters and deploy them with minimal effort.
- **GPU monitoring** – The ability to view GPU resource utilization across clusters and hosts alongside the existing host memory and capacity consoles enables admins to optimize GPU usage and make informed decisions.
- **PowerCLI scripts** – These powerful, customizable scripts automate the deployment of the infrastructure prerequisites when implementing AI workloads on VCF.
- **NVIDIA AI Enterprise** – This end-to-end, cloud-native software platform accelerates the data science pipeline and streamlines development and deployment of production-grade AI applications. It includes the following features:
  - **NVIDIA NIM** – A set of easy-to-use microservices designed to speed up the deployment of GenAI across the enterprise.
  - **NVIDIA NeMo Retriever** – A collection of NVIDIA CUDA-X GenAI microservices enabling your organization to seamlessly connect custom models to diverse business data and deliver highly accurate responses.
  - **NVIDIA RAG LLM Operator** – This operator streamlines the deployment of RAG pipelines into production, no rewriting of code required.
  - **NVIDIA GPU Operator** – This operator automates the lifecycle management of the software required to use GPUs with Kubernetes and enables advanced functionality such as better GPU performance, utilization and telemetry.
- **Major server OEM support** – VMware Private AI Foundation with NVIDIA is supported by major server OEMs such as Dell, Lenovo and HPE, enabling your organization to leverage the platform's advanced AI tools on trusted, high-performance hardware from industry leaders.
- **Private AI Automation Services** – This advanced VCF Automation capability comes with a Catalog Setup Wizard that provides a quick-start experience, automated setup of Private AI services, and self-service provisioning of GPU-capable machines, including Machine Learning (ML) workloads and VKS GPU-enabled Kubernetes clusters.

- **Private AI and GPU monitoring with VCF Operations** – VCF Operations enables you to monitor GPU metrics like compute, memory, temperature and power consumption at a cluster and host level for ESX hosts having GPUs, helping you to assess the performance-related effects of the app on the GPU hardware and on the system as a whole.
- **Comprehensive data services management** – VMware Data Services Manager provides the underlying vector database (pgvector) management to support the Private AI use case.

Key benefits of using VCF to run traditional workloads

- Lower operational overhead
- Maximized value and cost savings
- Reduced unplanned downtime
- Future-ready infrastructure

Helpful resources

VCF website

VCF datasheet

Private Cloud Modernization Program solution brief

VCF TCO white paper and infographic

VCF technical documentation

VCF Getting Started Hands on Lab and HoL catalog

VCF blogs

Follow us on X

Follow us on LinkedIn

Watch latest videos on YouTube

Use case: Run and manage traditional workloads

VCF enables you to run and manage traditional workloads such as ERP systems, CRM systems, databases, VDI workloads, financial apps, inventory management or supply chain apps, healthcare apps, and HR/Payroll apps without any downtime or security breach. Key features and capabilities include the following:

- **A comprehensive private-cloud platform** with integrated, enterprise-class compute, networking, storage, management, and security that can be deployed within hours instead of months or years.
- **Secure and resilient infrastructure for production workloads** – VCF provides robust security features such as microsegmentation, data-at-rest and data-in-transit encryption, identity and access management (RBAC), data governance, constraints-based management for compliance, advanced threat detection, and regulatory compliance enforcement. In addition, features such as vSAN Data Protection, Stretched Clusters, disaster recovery and avoidance, and ransomware protection and recovery help ensure minimal service disruption and data loss during unforeseen events.
- **Automated infrastructure** – With VCF Automation integrated into the platform, VCF delivers comprehensive management and ongoing maintenance capabilities for traditional workloads. With IaC, blueprints and a GitOps-based approach, it provides consistent configurations across environments, reducing discrepancies and increasing standardization for more reliable systems and minimal risk of configuration drift. In addition, it enforces policies and compliance requirements consistently across systems, reducing the risk of human error in configuration.
- **Continuous monitoring** – VCF Operations, another integral component of VCF, monitors mission-critical apps for performance, availability, and end-user experience with capabilities such as service discovery, agent-based monitoring (curated and open-source Telegraf agents), and third-party integrations with existing application performance management (APM) systems.
- **Intelligent operations** – VCF Operations delivers comprehensive operational capabilities such as performance optimization, capacity management, cost efficiency management, visibility and observability of infrastructure and operations, and compliance, troubleshooting, and log analytics.
- **Comprehensive data services management** – VMware Data Services Manager simplifies the deployment and management of open-source and commercial databases from a single pane of glass, and enables IT to offer self-service, enterprise-hardened Postgres, MySQL, and Google AlloyDB Omni (tech preview) to their teams, boosting developer innovation, reducing IT costs, and ensuring data resilience.

Get started

Learn more about how VCF can help you deliver a unified cloud experience. Want help in your cloud journey? Our Private Cloud Modernization Program is designed to guide you through every step, no matter where you are in the process. Please contact your Broadcom representative to learn more.

vmware® by Broadcom

Copyright © 2025 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.
Item No: FY25-7859-BC-VMW-VCF-UNIFIED-CLOUD-EXP-SO-USLET-WEB-20250409 4/25