Technical White Paper: **2023**

VMware Cloud Well-Architected Framework

vmware®

Table of contents

Executive summary

There are several 'well-architected' frameworks that exist in the market today. This guidance has primarily been focused on individual public cloud providers with an emphasis on cloud-only deployments. In contrast, VMware's customer environments span across different deployment models such as private cloud (on-premises), public cloud and hybrid cloud. As organizations increasingly adopt multiple cloud providers to support their business interests, a need has emerged to offer prescriptive guidance and best practices supporting multi-cloud deployments.

VMware is best positioned to deliver comprehensive guidance to support customers throughout their multi-cloud journey. This paper introduces the new VMware Cloud™ Well-Architected Framework (VMCWAF) that builds upon the following VMware pillars — Plan, Build, Secure, Modernize and Operate. These pillars will help customers execute against their multi-cloud strategy and will help organizations build reliable, scalable, secure and operationally efficient cloud environments.

For any feedback, suggestions or corrections, please send an email to: vmwcloudready@vmware.com.

Audience

This document provides an overview of a new VMware framework that aims to educate and guide customers throughout their multicloud journey. The intended audience for this whitepaper is technology executives (CTO/CIO), solution and cloud architects. These individuals are responsible for the technical direction and design of cloud infrastructure to support their organization's digital transformation initiatives.

Multi-cloud concepts

Organizations must decide on a cloud strategy that meets the needs of their business, whether that is hybrid cloud and/or multi-cloud. Hybrid cloud is defined as the use of private cloud and public cloud platforms to provide a flexible mix of cloud computing services, allowing for consistent infrastructure, simplified workload migration and placement. Multi-cloud, by comparison, is defined as the use of two or more public cloud providers with or without any existing private cloud infrastructure. Multi-cloud is an emerging strategy employed by organizations that need to meet specific technical requirements and business outcomes by leveraging services across multiple cloud providers simultaneously. Organizations adopt this approach when faced with a scenario where no one cloud platform is able to meet all their technical requirements or deliver all the necessary outcomes for the business. This often occurs through mergers and acquisitions (M&A), where businesses find themselves with application, operational and financial ownership across clouds that may not align with their initial cloud strategy. This model affords organizations the flexibility, choice and unique opportunity to be intentional about their approach to infrastructure and application modernization (e.g., refactor/build, replatform, rehost/migrate, retain or retire).

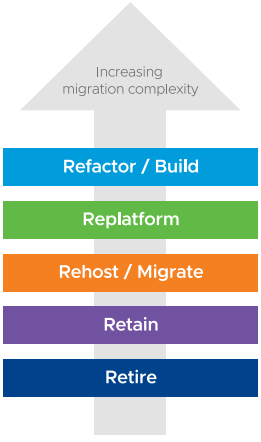

Figure 1: Common migration strategies

- **Refactor/Build** involves changing the application at the source code level. Typically, applications are re-written to take advantage of cloud micro-services architecture and to incorporate new services such as IoT, machine learning and others
- **Replatform** involves changing the operating system, such as going from Windows to Linux, modifying the application middleware, such as going from a self-managed database to a cloud provider managed database or from a virtual machine to a container image
- **Rehost/Migrate** involves either changing the hypervisor. (e.g., migrate applications from one virtualized environment to another) which is known as **Rehost** or moving an application without changing the underlying hypervisor or application at a source code level (e.g., migrate VMs from one virtualized environment to another without requiring changes) which is known as **Relocate**
- **Retain** means leaving workloads in a private cloud environment
- **Retire** means decommissioning workloads and/or converting to SaaS

For most organizations, the need for a multi-cloud strategy is rapidly becoming an inevitability. Considerations for multi-cloud are not limited to just technical decisions. Service level agreements (SLAs) and service level objectives (SLOs), instance types or node specifications, region/availability zone (AZ) placement, storage and network performance characteristics are important factors, but do not stand alone. Business considerations such as existing commercial agreements, regulatory compliance requirements, organizational culture and cloud perception, skills availability and many others are also critical concepts to weigh in when building out a complete multi-cloud strategy. The VMware Cloud Well-Architected Framework will introduce concepts, principles and best practices for navigating these scenarios as organizations move towards multi-cloud deployments.

Common drivers for multi-cloud environments

The challenge that organizations face today is not whether to adopt cloud, but how to map their requirements to the most suitable cloud provider for both current and future needs. Whether technical in nature or business-driven, considerations for multi-cloud can vary by organization as well as individual departments and business units. Considerations such as market competition, licensing agreements, data privacy, regulatory compliance and existing partnerships often introduce competing needs and bespoke views, heavily informing cloud adoption decisions.

To navigate this complexity with minimal trade-offs while also avoiding a 'lowest common denominator' concession scenario, organizations are looking to a multi-cloud model to leverage the differentiated native cloud services that are best suited to meet the needs and goals of the business. One such pattern is running certain workloads on a particular cloud to take advantage of extended or enhanced support capabilities — or avoid licensing cost increases by running elsewhere.

Customers also have a need to ensure that regulatory data-resiliency requirements can be achieved. Depending upon the geographical location and data requirements, this could create a situation where an organization's cloud provider of choice does not have sufficient regional or geographic presence to meet this need. This framework will consider drivers that include, but are not limited to, the following:

- Business agility and operational consistency
- Regulatory compliance alignment
- Application modernization and native services integration
- Cost, consumption and licensing models
- Reduced reliance on a single-provider
- Enhanced service availability

Multi-cloud use cases

Early on, cloud adoption was driven by an organization's technology adoption, which closely aligns with their culture (e.g., innovators, early or later adopters). The use cases for cloud computing have remained relatively consistent. Whether the goal is enhancing and/or extending disaster recovery (DR) capabilities, allowing for seamless rapid capacity increases, improving developer experience or reducing operational overhead, the strategic importance of each decision is defined solely by the value to the business.

As the industry has matured and centered around cloud, multi-cloud has emerged as the next phase of adoption and strategic importance for many organizations. This emergence holds equally true for companies that are still young in the cloud and addressing their adoption and migration strategies, as it does for those businesses already firmly planted in the execution of their cloud strategy, driving towards multi-cloud deployments. Below are a few examples that will be expanded upon in future iterations of the VMware Cloud Well-Architected Framework:

- Access to differentiated cloud provider native services
- Capacity augmentation and elasticity
- Datacenter evacuation and/or consolidation
- Geographic expansion
- Prescriptive application modernization
- Disaster recovery and business continuity
- Data sovereignty and regulatory compliance

VMware Cloud

VMware Cloud delivers a modern multi-cloud platform that provides unified infrastructure, management and operations and cloud services in order to help customers build and deploy modern applications, from the data center to multiple cloud and edge environments. It enables the ability to build, run, manage, connect and protect any application on any cloud. Customers can choose the best cloud or clouds, whether private, public, or hybrid, for their applications without re-architecting and maintaining the highest level of consistency for infrastructure, operations and experience.

With unified infrastructure, VMware Cloud provides the industry's best compute (vSphere®), storage (VMware vSAN™) and networking (VMware NSX®), integrated through VMware Cloud Foundation™, across any public cloud or private cloud hardware. With VMware Cloud's consistent infrastructure, customers can migrate workloads seamlessly between environments and ensure that all data and applications remain secure and protected in any cloud. With unified management and operations, VMware Cloud delivers self-service, automation and governance, performance, troubleshooting, capacity and cost analysis services that help customers accelerate innovation, gain efficiency, improve control and mitigate risks. And finally, with unified cloud services, VMware Cloud delivers support for accelerated migration services, disaster recovery solutions, 300+ ISV ecosystem solutions and advanced native cloud services from individual cloud providers.

VMware Cloud solutions span multiple hyperscale public cloud providers including Microsoft Azure, Google Cloud, IBM Cloud, Oracle Cloud and Alibaba Cloud, as well as our partnership with VMware Cloud on AWS — VMware's preferred public cloud partner for all vSphere-based workloads. With more and more customers embracing multi-cloud, this document complements and provides continuity with these providers' own architectural guidance and frameworks, referenced below. Apart from the large hyperscaler cloud partners, VMware Cloud offerings are available through our network of 4,300+ VMware Cloud Providers, including over 230 VMware Cloud Verified partners. Learn more about VMware Cloud.

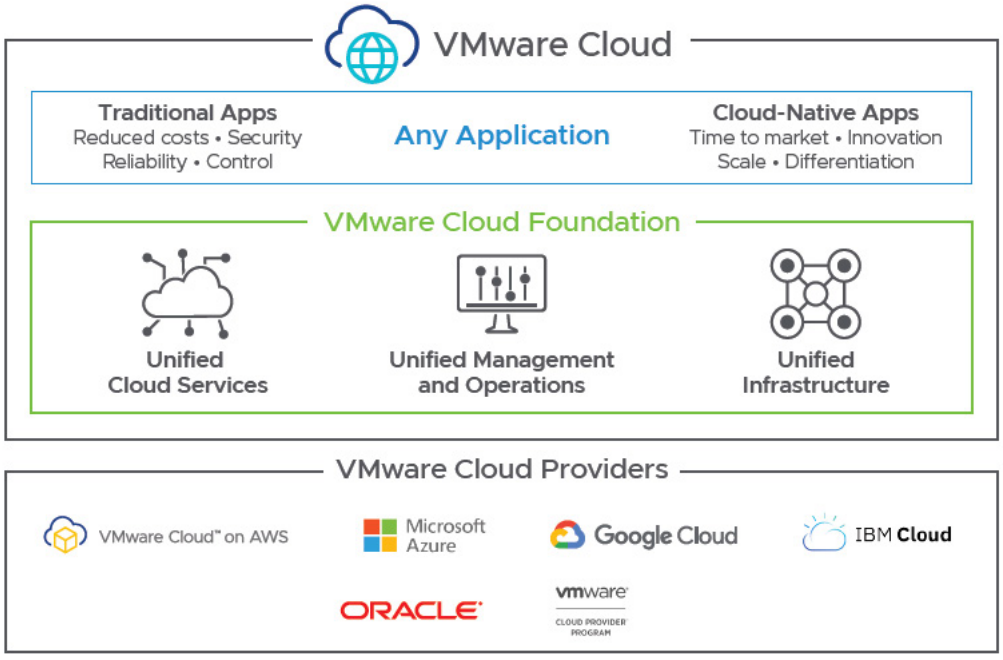

Figure 2: VMware Cloud platform

VMware Cloud™ on AWS

VMware Cloud on AWS brings VMware's enterprise class Software-Defined Data Center software to the AWS Cloud and enables customers to run production applications across vSphere-based private, public and hybrid cloud environments, with optimized access to AWS services. It integrates VMware's flagship compute, storage and network virtualization products (vSphere, vSAN and NSX) along with VMware vCenter® management, as well as robust disaster protection and optimizes it to run on dedicated, elastic, Amazon EC2 bare-metal infrastructure that is fully integrated as part of the AWS Cloud.

vmware | Microsoft Azure

Azure VMware Solution is a first party Microsoft Azure service that empowers customers to seamlessly extend or migrate existing on-premises VMware workloads to Azure without the cost, effort or risk of re-architecting or retooling. As a VMware Cloud Verified solution, you can keep using existing VMware investments, skills and tools you already know including vSphere, vSAN, NSX and vCenter while you modernize your applications with Azure native services.

vmware | ORACLE

Oracle Cloud VMware Solution provides customers a managed, native VMware-based cloud environment, installed within a customer's tenancy. It offers complete control using familiar VMware tools. Move or extend VMware-based workloads to the cloud without re-architecting applications or retooling operations. Oracle Cloud VMware Solution accelerates enterprise cloud migration with the ability to move seamlessly between on-premises and cloud environments, as well as hybrid deployments that span both.

vmware | Google Cloud

Google Cloud VMware Engine is a subscription cloud service developed and operated by Google and leverages VMware Cloud Foundation. The service is Cloud Verified and includes vSphere, vSAN and NSX deployed in Google's cloud platform and operated by Google. This allows for seamless migrations from on-premises environments to Google Cloud VMware Engine and the ability to integrate with Google Cloud services.

vmware | IBM Cloud

IBM Cloud® for VMware Solutions makes it simpler for your organization to capitalize on the tremendous potential of the cloud. Migrate VMware workloads to the IBM Cloud while using existing tools, technologies and skills from your on-premises environment. The integration and automation with a developer ready platform helps accelerate innovation with services like AI, analytics and more.

vmware CLOUD PROVIDER™ PROGRAM

VMware Cloud Providers™, part of the VMware Partner Connect program, include over 4,300 cloud providers operating in 120+ countries and powering over 10M workloads. As the world's largest network of validated cloud services based on VMware technology, these partners offer a wide variety of cloud services that includes disaster recovery, migration, container and application services — all built on the VMware Cloud Provider Platform. This diverse set of cloud providers also includes over 230 VMware Cloud Verified partners ready to bolster your business with the world's leading cloud infrastructure delivered as a service.

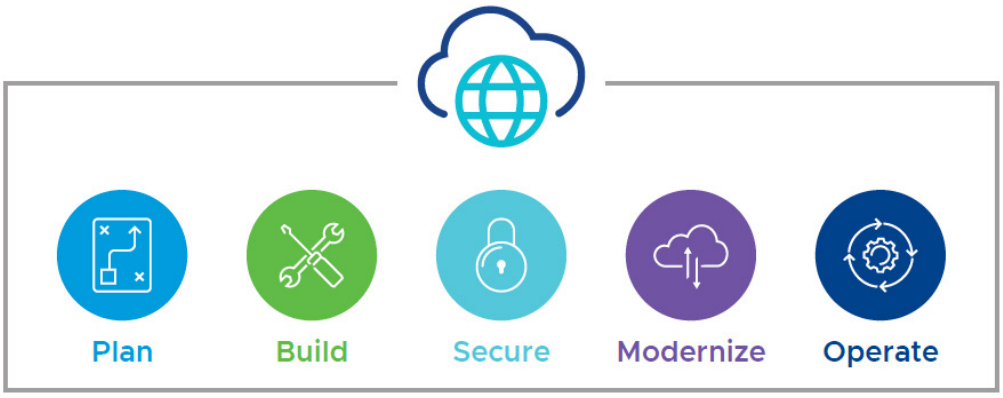

Figure 3: VMware Well-Architected Framework pillars

VMware pillars

Plan

The Plan pillar focuses on defining the scope and requirements as organizations transition to a cloud operating model that includes both existing and new applications. It is important to also review team structure, culture and skillsets prior to execution. This will highlight organizational optimizations that will lay the groundwork for transitioning to a multi-cloud model. Similarly, existing infrastructure and workloads should be assessed to identify their dependencies and the best operating location for each workload.

The initial assessment may include application and infrastructure dependencies, support levels, scale, high availability and sizing requirements prior to moving to the Build pillar. Observability and automation requirements should also be factored into the design and implementation as part of this process.

Concepts

- Organizational principles and culture
- Assessing existing workloads and infrastructure including application dependencies
- Designing the cloud environment for scale, high availability and recoverability
- Plan for Service Level Agreements (SLAs) and Service Level Objectives (SLOs)
- Provider SLA (customer managed, VMware managed, partner managed)
- Infrastructure/application
- Observability and system health
- Services (Directory Services, NTP, DNS, DHCP, Syslog, etc.)

Build

The Build pillar is where a VMware Cloud Software Defined Data Center (SDDC) is deployed and configured based on the information gathered from the Plan pillar. Next, it is important for an organization to implement supporting infrastructure services (e.g., NTP, DNS, DHCP, Syslog) to ensure a highly available and resilient environment. Once all infrastructure services are in place, organizations can configure identity management to secure user access and provide role-based controls. Automation of infrastructure, workload and security services is an important consideration in the deployment process for manageability, consistency and scalability. Finally, implement backup and disaster recovery services based on the requirements of each workload to ensure proper recoverability.

Concepts

- Identity and access management services
- Infrastructure services (e.g., NTP, DNS, DHCP, Syslog)
- Automation of infrastructure, workload and security services
- Backup and disaster recovery services
- Best practices for sustainability and carbon reduction

Secure

The Secure pillar provides guidelines and recommendations on how organizations can protect and secure workloads and applications in VMware Cloud. A shared responsibility model is used for security, which is implemented through the partnership between customers, VMware and the respective VMware Cloud provider. The underlying infrastructure (compute, network and storage) of a VMware Cloud based solution is managed and protected by VMware and/or the VMware Cloud provider.

An organization is responsible for securing and protecting their workloads and applications. This can be implemented by leveraging VMware's best-in-class security capabilities across compute, network and storage products that comprise the foundation of VMware Cloud. This security foundation provides visibility and distributed security controls to deliver comprehensive threat protection for users, applications and data. Organizations can implement granular access control through segmentation based on applications and workloads, while inspecting all traffic for anomalous behavior and attacks while streamlining security operations across their enterprise.

Concepts

- Identity and access management
  - Role based access control
  - External authentication (AD/SSO)
  - Multi-factor authentication
- Security and compliance frameworks
  - Geographical certifications
  - Governance certifications
  - Industry specific certifications
- Security controls and policies (compute, network, storage and application)
  - Network and micro-segmentation
  - Data in flight and data at rest encryption
  - Automated security and hardening
  - Manageability and traceability (audit logging)

Modernize

The Modernize pillar includes strategies for both infrastructure and application modernization. Infrastructure modernization is comprised of migrating existing workloads from a private cloud to a VMware-based Cloud. Application Modernization can consist of refactoring an existing application or building a new application.

Transforming an organization's infrastructure and/or applications is a continuous journey which can be broken down into incremental workstreams. The speed in which a modernization project is executed will largely depend on an organization's business outcomes and timelines. For example, to support more frequent changes, security best practices and novel uses of data, applications may need to follow modern application patterns, such as microservices architectures and API-first design. An organization may choose to modernize a select set of applications that fundamentally differentiate their business. Organizations also have the opportunity to optimize infrastructure operations, shift from capex to opex spending and access elastic capacity with minimal friction. To minimize these costs, an organization may decide to reduce their private cloud footprint by migrating workloads to their preferred public cloud provider.

It is important to understand that application modernization is not one specific approach but can be a combination of approaches.

One pattern could consist of refactoring your application with a complete or partial rewrite to innovate at the scale and speed of the business. An alternative pattern can also be hybrid, where a subset of an application's services are modernized by taking advantage of cloud native services while other remaining services are left unmodified.

Concepts

- Leverage automation for workload migration/deployment, configuration and management
- Considerations for running workloads that will support your organization's sustainability efforts
- Deliver on your app modernization strategy
  - Evaluate application against common migration strategies (5 R's)
  - Define future application state and create a transition plan for execution
- Validation of migrated and/or modernized workloads
  - Verify application functionality, performance and business value
- Apply workload best practices and optimizations

Operate

The Operate pillar provides guidance for how an organization should think about ongoing workload management and cloud operations. It is important to understand that a shared responsibility model is used when consuming a VMware Cloud-based offering between customers and the respective VMware Cloud provider. The management, health and lifecycle of the underlying hardware infrastructure (compute, network and storage) is the responsibility of the VMware Cloud Provider. This enables organizations to completely focus on the configuration, availability, observability and the recoverability of their applications. Organizations must shift their mindsets and develop new guidelines specifically for monitoring, licensing, security, logging, backups, disaster recovery and performance for applications.

Concepts

- Define Service Level Agreements (SLAs) and Service Level Objectives (SLOs)
- Determine observability and health metrics
  - Metrics, logs, events and traces
- Recoverability
  - Backup, disaster avoidance and disaster recovery
  - Incident response and mitigate operational risks
- Continuous assessment and optimization
  - Workload right sizing
  - Analysis for potential app modernization
  - Capacity planning
- Monitor and fine-tune organization's sustainability metrics

Summary

With the introduction of the VMware Cloud Well-Architected Framework, technology executives, solution and cloud architects will be able to leverage the prescriptive guidance and best practices for executing against a multi-cloud strategy.

As VMware continues to evolve and mature this framework in collaboration with our customers and partners, organizations will be empowered to accelerate their multi-cloud journey by using prescriptive guidance and best practices. The long-term vision for this framework will include detailed technical assets that will focus on the pillars that are defined in this framework (Plan, Build, Secure, Modernize and Operate).

With this new framework in place, it provides customers an ability to add new and unique digital experiences and integrations across VMware's Cloud Services. To complement the VMware Pillars, a rich set of tools and self-service capabilities can be developed with the help and feedback of customers and partners over time. If you are interested in participating and contributing as a design partner, please email: vmwcloudready@vmware.com.

Additional resources

- VMware Cloud Well-Architected Framework
- Architecting Your Multi-Cloud Environment
- VMware Cloud on AWS Shared Responsibility Model
- Application Modernization in a Multi-Cloud World
- Delivering Intrinsic Availability and Resiliency
- VMware Multi-Cloud Architecture Solutions
- Optimizing Multi-Cloud Cost
- Multi-Cloud Architect Blogs

Changelog

Initial Release	02/24/21
Removal of VMware Cloud on Dell EMC	Updated on 04/25/2023

Authors

William Lam, VMware

William is a Senior Staff Solution Architect working in the VMware Cloud team within the Cloud Services Business Unit (CSBU) at VMware. He focuses on Automation, Integration and Operation for the VMware Cloud Software Defined Datacenters (SDDC). One of his primary responsibility is driving VMware Cloud's Customer initiative and helping provide early feedback on the usability, design and architecture of new VMware Cloud features and capabilities.

Nick Marshall, VMware

Nick Marshall is a Staff Technical Product Manager working with the VMware Cloud team and helping lead the VMware Cloud Well-Architected program. He focusses on helping customers build out their multi-cloud architecture with reference designs and blueprints along with helping VMware partners integrate their offerings with VMware Cloud.

Adam Osterholt, VMware

Adam is Chief Technologist for Cloud Services and a member of the Global Field and Industry Program through the VMware Office of the CTO. He works with customers to design multi-cloud architectures that simplify the transition to public cloud and simplify migration and modernization of their workloads. Finally, he provides technical guidance on incubation of new cloud technologies across our portfolio of partners.

Emad Younis, VMware

Emad Younis is a Staff Cloud Solutions Architect in the Cloud Services Business Unit (CSBU) at VMware. He leads VMware's Multi-Cloud Center of Excellence, focusing on solving customer problems related to cloud architecture, connectivity, workload mobility and native services across all VMware-based clouds. His responsibilities include generating technical content, evangelism and working with VMware's hyperscaler partners to drive VMware Cloud adaption.

Item No: VMware Cloud Well-Architected Framework 4/23