



VMware Cloud Well-Architected Framework for Azure VMware Solution: Planning Principles

VMware Architecture

Table of contents

VMware Cloud Well-Architected Framework for Azure VMware Solution: Planning Principles	3
Organizational Principles and Culture	3
Communication	3
Executive Sponsorship and Alignment	3
Organizational Readiness	4
Service Level Agreements and Objectives	5
Assessing Existing Workloads and Infrastructure	6
Discovery	6
Infrastructure Inventory	6
Data Collection	6
Analysis	7
Requirements Gathering	7
VMware Cloud Migration Strategy	7
Migration Network Connectivity	8
Migration Wave Planning	9
VMware Cloud Design	9
Designing for Scale, High Availability, and Recoverability	11
Design Process	11
Scalability	11
Designing for Scale	11
High Availability	11
Designing for High Availability	12
Recoverability	12
Designing for Recoverability	12

VMware Cloud Well-Architected Framework for Azure VMware Solution: Planning Principles

Organizational Principles and Culture

Whether you call it Digital Transformation, Cloud First or Application Modernization, most businesses are going through some form of transformation to become more competitive in this new digital era. According to a [recent survey by VMware](#), 91% of executives agree their major application initiative in 2021 is to migrate and modernize legacy applications. The promise of cloud is its potential in transforming an organization to create and deliver new digital experiences for their end users.

Consumption of cloud services and new tooling will demand organizations learn new skills, adapt existing processes, and update automation workflows. The speed in which an organization can achieve cloud transformation will heavily depend on executive sponsorship, organizational readiness, and clear communications.

A traditional operating model relies predominately on managing the procurement of physical assets, lifecycle management, end-to-end visibility, and ownership across the entire technology stack. Transitioning to a cloud operating model will change the organizational dynamics that many organizations are currently operating under.

Communication

Cultural transformation starts with a clear purpose, supported by succinct and continuous communication. It is imperative that as organizations embark on the necessary changes to support their move to a cloud operating model, that the reasoning for the move is clearly understood. Document the business outcomes and ensure it is visible and available to all stakeholders in the organization.

To change the culture and operating model of an organization takes time. Prioritizing key workstreams, attaching timelines, and assigning clear ownership and accountability is critical to the transformation of the organization. Involvement from key stakeholders across the business is important in ensuring all requirements are communicated and captured early. Ambiguity will cause delays, budget overages, and may negatively impact business outcomes.

How an organization communicates can either inhibit or accelerate the desired culture and business outcomes. While not an exhaustive list, it is important to consider the following as an organization plans their communications to support a cloud operating model:

- **Communication Promotes Motivation**
 - Clearly articulate expectations and outcomes
 - Communicate the importance of each contribution against business outcomes
- **Ownership Drives Accountability**
 - Provide clear and consistent lines of accountability across areas of change
 - Provide a clear and documented scope of responsibilities
 - Provide clear guidance, measurements, and criteria for success
- **Inclusion Fosters Success**
 - Ensure those impacted by the transformation have their perspectives included
 - Create an environment for immediate feedback and contrarian views

With any organizational change, it is not uncommon for new requirements to surface, introducing unknown constraints during execution. Organizations should have processes in place to incorporate new requirements into existing workstreams.

Executive Sponsorship and Alignment

Digital transformation and Cloud Initiatives are most often born and driven by executive leadership. At the highest level of any organization, there is often strong alignment on these strategic initiatives, or at a minimum, agreement to proceed with organizational transformation. As organizations prepare for and execute against their transformational initiatives, issues will arise. It is imperative that support, alignment, and clear sponsorship at the executive level are present and accessible to those driving the supporting workstreams and keep the following considerations in mind:

- Hold themselves publicly accountable for its success
- Champion and lead by the example the cultural changes taking place

- Make timely decisions to positively impact progress
- Remove roadblocks to ensure consistent forward progress
- Allocate resources (people, infrastructure, tooling, licensing, etc.) as needed
- Remove stalemates to keep momentum and remove tension

As an organization prepares to pursue a transformative initiative, alignment and sponsorship will have a direct and measurable impact. It is important to continually have open discussions to ensure the vision, strategy and execution have alignment across stakeholders. Consider the following points as you prepare for the transformation:

- How do organizational stakeholders envision their overall, future Cloud position?
 - Are there broad understanding of what success looks?
- Who will lead the organization through its strategic planning including goals, objectives, and actions?
 - How are decisions communicated? Is feedback welcomed and open?
- How are success criteria for the cloud journey defined, and what are the specific and measurable KPIs to action against?
 - What are the measurements of success as defined by leadership, and do they correctly translate into KPIs that can be measured and managed?
- How are strategic management processes documented, regularly analyzed, and improved?
 - Is this information gated, or made openly available to all individuals in the organization?
- How does the current culture, and go-forward strategy align for long-term success of the organization during, and post transformation?
 - Does the organizational culture foster long-term transformational success?

Organizational Readiness

The ability for an organization to transform its culture is ultimately driven by its people. Processes must continually be adapted to meet new requirements, and the use of technology must continue evolve to meet the needs of the business, but the people in an organization is the constant in any transformation. Individual stakeholders must be empowered to drive and lead transformation across an organization. Below are example questions to review internally to help gauge the readiness of the organization on executing against its transformation initiative:

- Is there complete alignment across the organization on which infrastructure service provider(s) the business will align?
 - If not, how will this risk be mitigated to reduce the likelihood of delays to progress, rogue cloud deployments and associated costs?
- Do existing agreements exist with the providers of choice, or do they need to be negotiated?
 - Does the individual expertise exist internally to properly negotiate said agreements?
- Are finance and procurement teams aligned on how to move from a Capital Expenditure (CapEx) model, to one driven predominately by Operational Expenditures (OpEx) for infrastructure resource consumption?
 - Do budget allocations and structure easily allow for this shift?
- Do key individuals possess the technical skills across operations teams to support the businesses critical workloads and applications that are moving to, or will be developed in the cloud?
 - If so, how quickly can IT change their runbooks, their tooling, and extend their automation to factor in remote and unique infrastructure components?
 - If not, how do these individuals acquire the necessary skills - Retrain individuals? Hire desired skills into existing teams? Outsource?
- What is the internal perception of public cloud, and is it the same across all facets of the business?
 - Are there cloud adverse stakeholders as acting members of the decision tree?
- Does all existing systems and tooling have sufficient licensing to be deployed and run on remote, cloud provider owned infrastructure?
 - If not, what is involved in ensuring the proper cloud licensing is acquired?

The answers to these questions and many others will provide an organization an idea of the potential tradeoffs and/or risks. Organizations that engage in an open and early dialogue will gain a better understanding of the required changes to successfully execute their transformation.

Service Level Agreements and Objectives

One of the key differences between an on-premises environment and a VMware Cloud SDDC is in the responsibility of the infrastructure management. In an on-premises environment, an organization is responsible for managing the physical infrastructure, virtual infrastructure, and workloads. With VMware Cloud, an organization is primarily responsible for managing and operating their workloads, while the VMware Cloud Infrastructure Service provider manages the physical and virtual infrastructure.

Note: For more details, please refer to the VMware Cloud shared responsibility model.

It is important for an organization to understand the Service Level Agreements (SLAs) and Service Level Objectives (SLOs) for a given VMware Cloud Infrastructure Service provider to ensure it satisfies the needs of the organization. In VMware Cloud, an SLO defines the quality of service that a VMware Cloud Infrastructure Service provider delivers to an organization. An SLA is a legally binding contract between a VMware Cloud Infrastructure Service provider and an organization with specific terms and conditions of the SLOs.

When analyzing existing workloads, an organization should agree on a set of SLOs with the respective application owners. An organization should design a VMware Cloud environment that meets the SLOs for their end-users, while aligning to the established SLAs of the VMware Cloud Infrastructure Service provider. In addition, an organization should monitor and log key metrics to ensure that a VMware Cloud Infrastructure Service provider is meeting their SLAs.

Note: VMware vRealize Operations Cloud and VMware vRealize Log Insight Cloud can be used to monitor the VMware Cloud SDDC and its workloads.

Assessing Existing Workloads and Infrastructure

The initial assessment of the existing infrastructure and workloads is critical to enable an organization to successfully onboard into a VMware Cloud SDDC.

The assessment consists of two main phases: discovery and analysis. Information about the existing infrastructure and workloads will be collected during discovery, and then the appropriate cloud migration strategy will be determined during the analysis phase.

Discovery

Infrastructure Inventory

The first step of discovery is to build an accurate inventory within an organization's infrastructure that includes, but are not limited to, the following:

- Physical workloads (i.e., bare-metal nodes, firewalls, networks/VLANs)
- Virtual workloads (i.e., virtual machines, containers)
- Applications
- Third-party tools and integrations
- The inventory should include the application owner information, which can be used to conduct a more detailed interview.

An accurate inventory will help build a successful migration strategy. A combination of tools and/or interviews can be leveraged to create and validate the inventory.

Note: Tools such as the vSphere Client, PowerCLI, vRealize Operations Manager, vRealize Operations Cloud, Internal Change Management Databases (CMDB), and other third-party solutions can be used to gather an inventory of both physical and virtual workloads.

Data Collection

Upon completing the inventory collection, the next step is to collect detailed information about each application. Before an organization can determine the migration strategy for a particular application, a comprehensive understanding of its functionality is required. This involves the following information to be collected and validated through a combination of tools and interviews:

- Business function and criticality
 - How important is this application to the business?
 - What is the business impact if this application is not functional for a certain period?
 - The business impact can be tangible (i.e., lost inventory, legal penalties, lost transaction revenue) or intangible (i.e., brand damage, decrease in stock value, loss of employees)
- Is there an established SLA/SLO from the respective line of business within the organization?
- Compute and storage capacity requirements and resource utilization
 - What are the minimum capacity requirements?
 - What is the expected capacity growth over the next x years?
 - What is the current resource utilization?
 - Are there certain days and times or months in a year where there are peak demands for resources?
- Performance requirements (compute, storage, network)
 - What is the current performance baseline for compute, storage, and network?
 - Are there specific performance requirements that need to be met? (i.e., minimum IOPS, number of concurrent connections, provisioning time, etc.)
- Ingress and egress traffic flows and network utilization
 - Which network ports are required for traffic?
 - What are the average and peak network utilizations?
 - Are there periodic spikes in network utilization due to scheduled events, such as backup?
- Service dependencies (i.e., application dependency, third-party integration)

- What is the current application architecture?
- Does the application depend on other services and/or workloads for functionality?
- How often does the application communicate over the network?
- Business continuity and disaster recovery requirements
 - Is there any single point of failures (SPOFs) for the application that should be mitigated?
 - What are the recovery time objective (RTO) and recovery point objective (RPO) requirements?

Note: Tools such as VMware vRealize Operations Manager, vRealize Operations Cloud and vRealize Network Insight, vRealize Network Insight Cloud can be used to collect this information. vRealize Operations can analyze the current resource consumption of virtual machines and recommend sizing for virtual machines, which can be used in cloud migration planning. vRealize Network Insight can provide data on traffic usage and help identify or validate firewall requirements. Other third-party solutions can also be used to collect similar information.

Analysis

Requirements Gathering

Information gathered during the inventory collection is used to build a list of requirements for each application. For each application, the requirements can be organized by design attributes, that includes, but is not limited to, the following:

- **Scalability:** the ability for a system to continue providing the same level of performance or functionality when there is a change in utilization.
 - For example, an application architect has a requirement that the underlying infrastructure must be able to scale out dynamically if there is an increase in traffic load. This may translate to a scalability requirement where the underlying infrastructure must be able to add or remove a compute capacity within an acceptable time frame.
- **Availability:** the ability for a system to continuously operate and function for an extended time without interruption.
 - For example, the Virtual Machine workloads has a requirement that the underlying virtual infrastructure can provide an SLA of 99% uptime per month for management. This may translate to an availability requirement of a VMware Cloud SDDC having a minimum of 99% uptime or greater.
- **Recoverability:** the ability for a system to recover from a disaster or a failure.
 - For example, an application architect requires the data for the application cannot tolerate data loss for more than 2hrs. This may translate to a recoverability requirement where the date for the Recovery Point Objective (RPO) must not exceed 2hrs.
- **Manageability:** the measure of how easily a system can be deployed, configured, and controlled.
 - For example, an organization has a requirement to provide end users with a solution that enables self-service workload provisioning with governance. This may translate to an evaluation of a cloud management platform (CMP) that integrates with a VMware Cloud SDDC.

Note: Tools such as VMware vRealize Automation Cloud and other third-party solutions that are supported with VMware Cloud can be used as a CMP.

- **Performance:** the measure of how well a system accomplishes a given task.
 - For example, the application has a requirement to deliver at least 100 concurrent requests per minute to satisfy its service SLA. This will translate to the supporting virtual infrastructure provisioned with the necessary compute, network, and storage resources to meet the application requirement.

The finalized list of requirements must be validated with the appropriate stakeholders within the organization through workshops or interviews before continuing with the assessment.

VMware Cloud Migration Strategy

With clear requirements and a holistic understanding of the application inventory, an organization can now determine an appropriate migration strategy for each application based on the needs of the business.

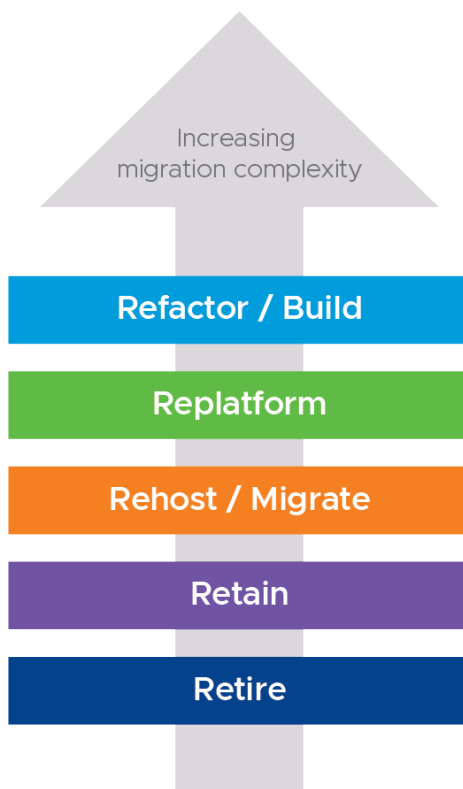


Figure 1: Common migration strategies

- Refactor / Build involves changing the application at the source code level. Typically, applications are re-written to take advantage of cloud microservices 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 and/or applications in a private cloud environment
- Retire means decommissioning workloads and/or applications, which can involve eliminating them altogether or converting to SaaS

It is important to understand that application modernization is not one specific approach but can be a combination of approaches. A common strategy that organizations have adopted to help accelerate their application modernization journey is first to migrate their existing workloads to a VMware Cloud SDDC and then modernize the underlying application.

Certain migration strategies, such as Refactor and Replatform, provide an opportunity for an organization to modernize their applications after migrating to VMware Cloud. The speed at which a modernization project is executed largely depends on an organization's business outcomes and timelines.

Organizations will be most successful in achieving their application modernization (app modernization) goals by leveraging a Migrate and Modernize strategy. By migrating existing Virtual Machine workloads to VMware Cloud, organizations will now have a modern infrastructure platform. Organizations will now be able to focus on their app modernization efforts.

Migration Network Connectivity

For each migration wave, the migration path and method must be determined.

There are various network connectivity options to create a migration path from an on-premises environment to a VMware Cloud SDDC, such as using a VPN or setting up a direct, private connection. Analyzing findings from the infrastructure assessment would

help identify feasible network connectivity options for an organization's specific business needs and requirements.

Note: VMware HCX can be used to provide a private, secure connection and migrate workloads between an on-premises environment and a VMware Cloud based SDDC. VMware HCX also provides different migration methods, such as cold migration, bulk migration, and live migration.

Migration Wave Planning

Wave planning is the process of grouping workloads that will be migrated concurrently based on business critically and application dependencies to help create a high-level migration schedule. Workloads can be migrated based on the application SLAs, for example non-mission critical workloads can be migrated initially.

It is also critical to understand the different types of migration methods and an organization should select the one based on the needs of the business. For example, a Dev/Test workload which can afford downtime during the evenings, a production workload that cannot afford any downtime, and a staging workload which can have minimal downtime when scheduled. From a migration execution standpoint, you would then select three different migration types as mentioned below for each of the respective workloads, maximizing the speed at which you can migrate the workloads and maintaining the application service level agreements (SLA).

There are different methods for migrating workloads such as hot, warm, and cold migration:

- A hot migration is referred to as a live migration and is the most familiar to VMware administrators. It is a staged migration where the virtual machine stays powered on during the initial full synchronization and the subsequent delta sync, using the VMware vSphere® vMotion® feature.
- A warm migration is a virtual machine that is actively running while it is being replicated to ensure minimal downtime. After the migration completes, you either start a manual or automated cutover to make the replicated virtual machine available on the cloud provider. Cutover is a process of powering on the virtual machines at the cloud provider site after the warm migration gets completed. This cutover operation includes a final sync and import of the migrated VM into a destination VMware Cloud SDDC.
- A cold migration is a virtual machine that is in a powered-off state before starting the migration. Exporting and importing virtual machine images is another form of cold migration.

Migration waves should also incorporate the related application dependencies and network communication traffic to keep intra-application traffic within the same environment and limit traffic across an on-premises data center and/or a VMware Cloud based SDDC.

In addition to grouping by applications, isolated waves should be created for large or complex workloads, such as database virtual machines or virtual machines with a high data change rate. These workloads tend to require more network bandwidth for migration and could impact other migrations if performed concurrently.

Note: VMware vRealize Network Insight and VMware vRealize Network Insight Cloud can be used to validate application dependencies and traffic flows. VMware HCX integrates with both VMware vRealize Network Insight and VMware vRealize Network Insight Cloud can automatically create a VMware HCX Mobility Group which migrates pre-defined set of workloads based on a migration wave planning.

VMware Cloud Design

After a thorough assessment of the existing infrastructure and workloads, the results will guide an organization in creating a VMware Cloud SDDC using design decisions based on compute and storage sizing, service location selection, and network connectivity.

The requirements gathered during the assessment will guide the design decisions for all aspects of a VMware Cloud SDDC, such as compute and storage sizing, service location selection, and network connectivity.

Appropriate compute and storage sizing must be determined for the VMware Cloud SDDC include resources for management components and overhead as well as the expected growth of workloads when sizing the VMware Cloud environment.

Note: [VMware Cloud Sizer](#) should be used to help an organization size the VMware Cloud based SDDC. The [VMware Cloud Configuration Maximums](#) document should also be referenced to ensure scalability of a VMware Cloud SDDC.

Service location for a VMware Cloud SDDC should be chosen depending on the requirements, such as the following:

- Availability of services: not all VMware Cloud services are available in every region
- User locations: depending on the business needs (i.e., market expansion, local security compliance) and application

requirements (i.e., service feature availability, minimum latency for optimal performance), an organization's service(s) may need to be in close proximity to their end users.

Network connectivity for a VMware Cloud SDDC depends on the business needs and requirements. If an organization decides to keep the on-premises environment and use VMware Cloud for bursting to meet unexpected demands or for disaster recovery of the on-premises environment, permanent network connectivity may be needed between the two environments. If an organization decides to decommission the on-premises environment, then only network connectivity for the migration will be needed. In addition to deciding on the longevity of a network connection, data on network utilization and application traffic flows collected during the assessment should be utilized to determine the type of network connectivity and the required bandwidth.

It is important to remember that a VMware Cloud SDDC design must meet all the identified requirements. A detailed example of how a VMware Cloud based SDDC can be designed to meet availability, recoverability, and scalability requirements is discussed in the next section: Designing for Scale, High Availability, and Recoverability.

Designing for Scale, High Availability, and Recoverability

Similar to an on-premises infrastructure, a VMware Cloud environment must also be designed for high availability, recoverability, and scalability. The specific technical configurations will vary based on the specific VMware Cloud Infrastructure Service providers. The design process and considerations discussed in the following sections apply to any VMware Cloud based environment.

Design Process

Requirements for the VMware Cloud SDDC are derived from the application requirements gathered during the initial assessment of existing workloads and infrastructure.

There are many ways to design each part of a VMware Cloud environment, such as host clusters and network connectivity. Regardless, the final design should meet identified requirements within constraints and assumptions. Constraints are any limiting factors that may affect the design. They can be project-related, such as budget or timeline, or technical, such as an existing application architecture that cannot be altered. Assumptions can be made during the discovery, as not all the necessary information may be available immediately. Identified assumptions must be validated to make design decisions based on correct information.

In addition, any risks, whether they are technical or non-technical, that may arise from deciding on a particular design should be documented and addressed with potential mitigation strategies. Managing risk can vary based on the impact and level of effort. High impact risk items should likely be mitigated, whereas the organization may accept low impact risk items.

Once a particular design is finalized, it is important to document the decision, including justification and any related risks. It is also valuable to note which requirements have been fulfilled by a particular design to ensure that the final design meets all the identified requirements and business needs.

Scalability

Scalability does not have a standard metric used across the industry. Generally, performance or load testing can be done on a system to determine how flexible it can handle changes in demand. For example, performance testing can be done to measure how long a system takes to add more resources and meet peak demand. A system that only takes one minute to add more resources is more scalable than a system that takes an hour to perform the same action.

Designing for Scale

There are several options for scalability of a VMware Cloud SDDC. An organization may begin with scaling up a vSphere cluster by adding hosts to meet an increase in demand. Organizations can enable automatic resource allocations such as Elastic Distributed Resource Scheduler (eDRS) which is available in a VMware Cloud on AWS SDDC. If an automated resource allocation service is not available in a VMware Cloud SDDC, the process of adding a host can be automated by leveraging the VMware Cloud SDDC APIs.

If scaling up a vSphere cluster reaches the VMware Cloud configuration maximums or does not fulfill the business needs, the next option is to scale out by creating additional vSphere clusters within the same VMware Cloud SDDC. It is important to determine when a new vSphere cluster should be created to plan for day two operations and manage the growth of an environment.

Depending on the business needs and the estimated growth, an organization may choose to create multiple VMware Cloud SDDCs instead of scaling up a single SDDC. It is essential to identify any workloads that must communicate with each other spanning several VMware Cloud SDDCs to design the network connectivity between these environments appropriately.

In addition to the virtual infrastructure, the VMware management Virtual Machines should also be considered when designing for scalability. Typically, the management Virtual Machines, such as the vCenter Server and NSX managers, will be deployed with a pre-defined set of compute and storage resources. Based on the workload requirements and expected utilization collected during the initial assessment, the VMware management Virtual Machines may need to be resized if this capability is available within a VMware Cloud SDDC.

Overall, a VMware Cloud environment should be designed to be repeatable. A modular design makes an environment easier to replicate or expand across different regions to meet fast-growing demand. An organization should plan how their VMware Cloud SDDC will expand when designing the first VMware Cloud SDDC to expedite the deployment of a new VMware Cloud SDDC as well as to simplify cloud management and day two operations in the future.

Note: VMware Cloud Sizer can be used to size the VMware Cloud based SDDC with data collected during the initial assessment. Configuration maximums document can also be referenced to ensure scalability of the VMware Cloud SDDC.

High Availability

High availability can be achieved for physical infrastructure, virtual infrastructure, and application services. High availability is measured by uptime, the amount of time that a service has been operational.

Designing for High Availability

The VMware Cloud Infrastructure Service providers typically have multiple physical data centers in various regions throughout the world. The data centers in each region are designed to be independent of one another so that a failure in one data center would not affect another. An organization can choose to deploy their VMware Cloud SDDC in one or more supported regions. The data centers in each region are designed to be independent of one another so that a failure in one data center would not affect another. The VMware Cloud Infrastructure Service provider is responsible for providing high availability for the physical infrastructure according to the Service Level Agreements (SLAs).

A VMware Cloud SDDC inherently provides high availability for the Virtual Machines running in the SDDC using vSphere High Availability (HA) and vSAN storage policies. When an ESXi host fails, vSphere HA automatically restarts the Virtual Machines from the failed ESXi host to other ESXi hosts within the same vSphere cluster. vSAN storage policies provide data redundancy through appropriate RAID configurations depending on the number of ESXi host failures an organization can tolerate.

Although a VMware Cloud SDDC provides native high availability capabilities, designing a virtual infrastructure that meets the applications SLAs is ultimately the responsibility of an organization.

An organization can deploy multiple independent VMware Cloud SDDCs, each in a different region where the VMware Cloud Infrastructure Services provider is available. Workloads can be deployed across these environments to satisfy availability requirements. With this method, it is vital to understand the application dependencies and network utilization to design an appropriate network connectivity between the different VMware Cloud SDDCs.

If a [VMware Stretched Cluster](#) is available, a VMware Cloud SDDC can span across multiple geographical regions. Deploying a VMware Stretched Cluster will incur additional cost due to cross-regional replication traffic. A VMware Stretched Cluster can achieve higher levels of availability. An additional benefit of deploying a VMware Stretched Cluster is the ability to provide an extra level of local site protection for Virtual Machines by distributing the placement of Virtual Machines across regions.

Recoverability

Recoverability is measured by two primary metrics:

- Recovery Point Objective (RPO): the amount of data loss an organization can tolerate.
- Recovery Time Objective (RTO): the amount of downtime an organization can tolerate.

RPO is a point in time where a VMware Cloud SDDC and an organization's Workloads and applications can be restored. RTO is the amount of time it takes to restore a VMware Cloud SDDC and an organization's Workloads and applications after a failure.

Designing for Recoverability

The VMware Cloud Infrastructure Service provider is responsible for the recoverability of the VMware Cloud SDDC management Virtual Machines. However, simply relying on the virtual infrastructure SLA may not be sufficient on meeting the application requirements. An organization must be prepared for individual Virtual Machine failures by designing a proper disaster recovery and backup solution.

For disaster recovery planning, applications should be grouped based on business criticality, RPO/RTO requirements, and application dependencies. The recovery process should prioritize mission-critical applications with lower RTOs. Inter-dependent applications should be recovered together to ensure proper functionality. It is important to choose an appropriate disaster recovery solution based on the application RPO requirements. To meet a lower RTO, an organization can automate the recovery process such as Virtual Machine failover and failback.

Note: Tools such as [VMware Cloud Disaster Recovery](#) and [VMware Site Recovery](#) can provide disaster recovery capabilities for a VMware Cloud SDDC.

In addition, proper monitoring must be in place to ensure that a Virtual Machine or an application failure is detected as soon as possible. The monitoring tool can be configured to alert on specific infrastructure and/or application issues and provides a means to notify the responsible recipients.

Workloads and/or application-level backups are critical to having a comprehensive disaster recovery solution. Backup retention policies and backup job scheduling should be configured to meet an organization's RPO requirements. Appropriate network connectivity with sufficient bandwidth should be provisioned to ensure backup jobs do not affect production traffic. A backup window should be scheduled outside of normal business hours to avoid impact to production workloads. Backups should be stored offsite, in a different location from where the workloads are residing in. An organization should regularly test and validate backups to ensure proper workload recoverability.

A VMware Cloud SDDC can also be used as a disaster recovery destination for an on-premises environment. An organization should have appropriate network connectivity to ensure their users can continue to access their workloads after a disaster has been

declared. A plan should be in place for an organization to fail back workloads to their original location once the infrastructure has been restored. An organization should regularly test and failover workloads to and from a VMware Cloud SDDC to validate their disaster recovery procedures.

In the next section, learn about managing and accessing costs for cloud infrastructure providers.

