White Paper: **November 2021**

VMware Cloud Director Availability

Migration White Paper

vmware®

Table of contents

Introduction to VMware Cloud Director Availability

VMware Cloud Director Availability is a powerful solution built to offer simple, secure, and cost-effective onboarding, migration, and disaster recovery services “to” or “between” multi-tenant VMware clouds. VCDA provides simplified self-service migration experience to end-customers.

Using VMware Cloud Director Availability, you can migrate your applications to a Cloud Provider without refactoring your application or doing any architecture changes, vSphere Virtual Machines are migrated ‘as is’ to the cloud with minimal or no changes.

On-Premises site doesn’t require any additional changes to deploy VCDA appliance and pair it to the Cloud Provider. On-premises appliance requires only outbound TLS (usually over port 443) connectivity to the Cloud site over the internet or private link if available. During pairing tenant can restrict access from the Cloud site which will prevent Cloud Director users and the provider to configure migrations.

VMware Cloud Director Availability provides asynchronous host-based replication (HBR). Virtual Machine snapshot are not used as part of the replication process unless quiescing is used. Windows or Linux quiescing is available only for virtual machines that support quiescing. For more information, see **Guest OS Quiescing Support**.

Migration use cases

There could be several reasons to decommission legacy IT infrastructure. Hardware is out of support, cannot utilize modern technologies, operational cost is very high are several examples why migration to the cloud could be a wise choice.

Expired hardware support

Many on-premises environments are running old hardware which is out of support or cost for purchasing extended support is very high. Also, not vendors offer extended support and they prefer to sell new equipment or licenses which may not be compatible with existing old equipment or legacy software licenses.

At the same time running business critical applications on infrastructure with no vendor support is extremely risky and company policies usually prohibit this. Any attempts to solve such situation are expensive and do not provide long term solution.

Limitation to use cutting edge technologies

Many organizations have in their on-premises old datacenters, they run many applications directly on the bare metal or use old virtualization software, not all of them own licenses that allow these virtualized infrastructures to operate in high availability mode. Majority of them rely on out-of-the-box features for monitoring and manageability which are far away from capabilities of dedicated products. From security perspective observations are even worse – only a small number of these datacenters utilize modern security technologies and rely on ad-hock activities which makes it almost impossible to provide structured and consistent security approach.

High operational cost

Utilizing legacy on-premises datacenter is usually one of the biggest expenses for one organization. Powering and cooling a datacenter cost a lot, buying extended support packs costs a lot, not all vendors offer extended support, finding compatible additional resources sometimes is not possible, many examples can be added that show operating a legacy datacenter is expensive and outcomes are not comparable. Finding personal could also be a challenge. IT people are looking at modern technologies and shift their skillset quickly to keep themselves prepared for today's workforce market demands.

Benefits from migration to the cloud

With migration to the cloud organizations can gain several benefits. First migration to the cloud will allow to better suite different application to proper infrastructure. Highly demanding applications could be run on more productive environment, could be placed on higher storage tier, adding additional compute capacity happens in the background – things that are not so easy to get into the legacy datacenter. Cloud providers utilize latest virtualization software with licenses that unlock all enterprise features and tenant organizations benefit of this transparently.

In addition, cloud providers may offer subscription to different services which do not exist in the legacy datacenter and tenant may purchase one or another service to add additional level of security, data protection, observability, etc.

Second benefit are the new capabilities open to tenant. Cloud providers are entitled through VMware Cloud Provider Program to all VMware products which enables them to use the latest technologies for security, monitoring, analytics. Migration from legacy datacenter to a cloud provider datacenter unlocks all these capabilities to the tenant organization. A huge feature set of latest NSX products is exposed through VCD tenant portal and enables tenant to secure applications and data to a level that never will be possible in on-premises datacenter. Many cloud providers grant their tenants with access to VROPs and VRNI portals and enable the tenant to monitor every compute resource metric and track network packet flow in tenant's virtual datacenter. All these details allow tenant to take an educated decision when fine tuning of compute resources for specific application is necessary or when network flow and security need to be optimized.

Modern cloud provider infrastructure opens the door to change the way how business applications have been developed and used in legacy IT infrastructure. Tenant may start planning to replace the old monolithic application with modern apps based on microservice architecture.

Third benefit is the way how expense profile changes. Entire legacy IT related CAPEX turns into OPEX. And this OPEX is not mandatory to high and flat. Usually cloud providers offer Pay-As-You-Go model for tenant's virtual datacenter which enables tenant to use flexible amount of resources and pay only for resources that have been utilized for the last billing period. Applications that have higher utilization only several times in the year will have different cost each month. Many cloud providers offer also managed services which removes the requirement to run internal IT team.

Types of Migration

Based on the source and destination type there are two types of migrations: On-Prem to Cloud and Cloud to Cloud. Each type of migration has its own specifics requirements and setup.

On-prem to Cloud

You can use VCDA to migrate your existing vSphere VMs to VMware Cloud Director based public cloud when you want to decommission your legacy environment or if you want to expand your existing infrastructure to the cloud.

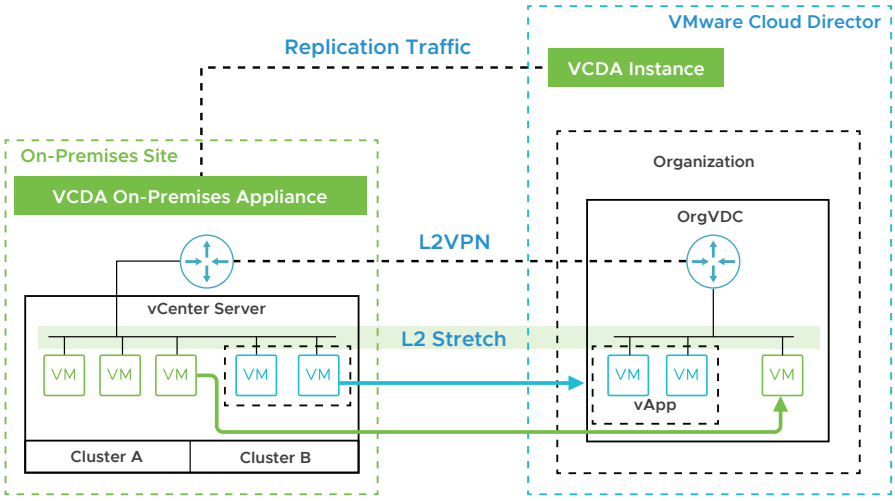

Figure 1: Example of On-prem to cloud migration setup

*L2VPN is optional and requires Autonomous NSX Edge to be installed and configured in the on-prem site.

The following table shows pairing Interoperability Between the Version of VMware Cloud Director Availability On-Premises Appliance the Version of the VMware Cloud Director Availability in the Cloud Site

Cloud On-prem	3.0	3.5	4.0	4.1	4.2	4.3
3.0	Supported	Supported	Supported	Unsupported	Unsupported	Unsupported
3.5	Supported	Supported	Supported	Supported	Unsupported	Unsupported
4.0	Supported	Supported	Supported	Supported	Supported	Unsupported
4.1	Unsupported	Supported	Supported	Supported	Supported	Supported
4.2	Unsupported	Unsupported	Supported	Supported	Supported	Supported
4.3	Unsupported	Unsupported	Unsupported	Supported	Supported	Supported

For information about the VMware Cloud Director Availability interoperability with other VMware products, see **VMware Product Interoperability Matrices**.

Cloud to Cloud

You can use VCDA to migrate existing workload VMs/vApps between Cloud Director based clouds. Such use case can be migration between different Cloud Director instances or Organization VCDs within the same or even different datacenters. It's useful when you want to migrate workloads from your legacy infrastructure. When you migrate vApps/VMs between Cloud sites information such as vApp/VM name, Guest properties and customization, description, boot order, vApp Networks and metadata is preserved.

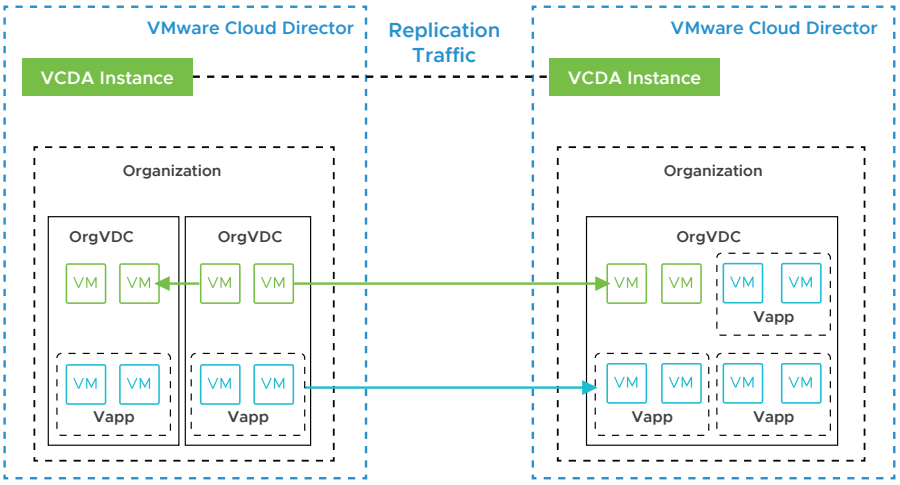

Figure 2: Example of Cloud to Cloud migration setup

Plan for migration

Migration preparation is important due to multiple reasons especially when downtime is not an option and not all workloads can be migrated in a short period of time. Preparation will include as minimum collecting information about:

- Grouping VMs by business application
- Storage, CPU and Memory consumption
- Network connectivity Identify Applications and group VMs

It is important to assess your current infrastructure and properly identify the virtual machines that will be migrated. For example, you might want to split your migration process into different maintenance windows for various reasons like business and process requirements, application types or resource constrains.

Identify the current CPU, Memory and Storage usage at your on-prem and plan your Cloud resources accordingly. If required increase your Organization VDCs resource to have enough resource for the migrated workload VMs.

Understand the difference between the VCD Allocation Models and how this can affect the migrated workload, for example CPU/Memory limits and reservations might be different between source and destinations sites.

For more information about VCD Allocation Models check this **link**.

It is crucial to plan the migration appropriate, for example you can chose to group the VMs based on their application. For example, you can have multi-tier application that spans multiple virtual machines VCDA give you the option to group those VMs in a single vApp at the destination Cloud side and to configure boot order and delay when performing the migration.

Grouping VMs

Many business applications are built out of several VMs so it's important to determine the groups of VMs representing different business applications. After determining which VMs build a single application it's required to analyze the communication flow as this will determine which networks from source will have to be extended to the destination. Another consideration that needs to be addressed is the intensity of communication between VMs and if they communicate in layer 2 broadcast domain or routing is involved. The plan for migration should account this information and migrations should be planned with reducing the number of communication streams over the intersite link between routed stretched networks and when this happens to try reducing the time for this state.

Grouping VMs by business application can happen by using analytics products like VRNI and VROPS. They can do such grouping by monitoring traffic between VMs in virtual infrastructure but confirmation by administrators is always recommended as traffic between different VMs may not be only application traffic but also management which in some cases may take a good portion out of total traffic to a VM (for example delivering OS updates to a VM). As many on-premises infrastructures do not have such instruments relying to application admin is the only way to determine which VMs are part of a single application.

Storage consumption

When migrating workloads from on-premises to the cloud storage consumption is not a huge challenge. The reason is VCDA will always create thin disks in destination no matter if disks in the source are in thin or thick format. In worst case scenario there will be no storage consumption reduction in destination compared to source site.

When migration between clouds is performed things are different and needs a special attention. VMware Cloud Director allows to enable an Organization VDCs for Fast Provisioning. If so, linked clones are used when creating vApps/VMs from templates which creates a significant reduction between provisioned and allocated storage space. A linked clone is a duplicate of a virtual machine that uses the same base disk as the original, with a chain of delta disks to track the differences between the original and the clone. As the migration cannot keep the same chain of base and delta disks in destination all disks will be with the size of a full clone no matter if Fast Provisioning is enabled for destination Organization VDC. This means every single VM disk will have the total size of its base disk plus all deltas before its own delta disk.

When a migration is configured VCDA will create a named disk in VCD to reserve space from organization's storage quota, this named disk will be equal to the provisioned size of all source disks, once initial synchronization is completed this named disk will shrink to the size of the thin disk (used space). When you perform the actual migration storage usage of a VM in VCD will equal to the provisioned space. You need to make sure that your storage policy quota has enough space for the migrated VMs, keep in mind that space for VM swap files is also taken from the storage quota.

As a result, the storage space reduction by Fast Provisioning won't be available in destination cloud and it must have enough space to accommodate all migrated workloads. Also, if migration was predeceased by a replication and this migration was configured to keep multiple point-in-time copies storage are destination site must be sized for capacity and performance to consolidate PITs in base image.

You can use **Cloud Director Availability Storage Calc** to plan the storage consumption and estimate the time required for sync to complete (when using the calculator please note that for "migration" RPO is 1440m and No. of MPITS is 1)

Network

Using VCDA 4.2 tenant may extend desired Layer 2 broadcast domains from its on-premises infrastructure to a cloud which is backed by NSX-T. If cloud provider uses NSX for vSphere tenant can do the same using VCDA 4.2.1. This allows extending migration windows without interrupting connectivity between VMs in source and destination.

This has its cost which must be considered when planning migrations.

All networks will be routed in source site. This means all ingress/egress Internet traffic will be via source gateway. For VMs that are migrated this means traffic will travel over the stretch which will result in increased latency. This latency may increase during times when VM disks are migrated between source and destination.

Routing between stretched network will happen on the router in source site. This means if VMs in different subnets are already migrated they will communicate over the inter-site link and router in source site.

For improving this situation, it's possible to re-configure routing through cloud earlier before migration of all workloads is completed. In this way inter-site link will be used for reduced VM-to-VM traffic over stretched networks and for replications and overall impact of migrations over services/applications will be minimized.

Planning your network setup is one of the most important steps when you are migrating your workloads especially when we talk about on-prem to cloud. You can choose VMs to keep the existing IP addresses or to change the IP addresses during the migration.

You need to prepare your organization networks in advance before the actual failover to the cloud side, during migration in VCDA you can select to which existing organization network to connect each Virtual Machine NIC.

If you decide to retain existing IP addresses and you cannot migrate all VMs that are in the same network (same L2 domain) during single maintenance window you can stretch your Layer 2 network to the cloud using VCDA. VCDA uses NSX and its L2 VPN services to stretch networks which allow network connectivity between migrated and non-migrated VMs. Migrated VMs on the stretched network will continue to use the existing routing in the on-prem side.

As of version 4.2.1 VCDA support both NSX-V and NSX-T in the cloud side for on-prem autonomous NSX Edge is used.

Migration

- **Migration configuration** - this is the step where you select your source VM/vApp for migration. At this point you select destination Organization, Org VDC and storage policy, destination storage policy must support "Named Disk" and "Virtual machines" entity types for successful configuration and migration. In case of on-prem to cloud migration of multiple VMs you have the option to group the VMs in a single vApp, configure boot order and delay. You can choose to delay the initial sync in case you have multiple syncs running at the same time and you have limited network bandwidth. Also, you can exclude some VM Disks or use seed VMs.
- **Initial sync** - once configuration is completed initial sync will start automatically if the VM is powered-on, if the VM is powered-off sync must be started manually. VCDA migrations sync automatically every 24h, if you have configured "Replication" you can chose the PRO interval.
- **Migrate/Failover Network Settings** – once the migration is configured you need to configure the network setting that will be applied once the VM is migrated, if not configured you can choose the destination network during the actual migration. You can configure both **Migration** and **Test** network settings. For **Test** setting you can use an isolated network to avoid any IP conflicts or network disruptions since when you do a test failover the source VM will remain Powered-on. You can also configure Guest customization and Computer name keep in mind that Guest Customization must be enabled if you want to change the VM IP address or computer name.
- **Test Failover** – it's always good to do a test failover before the actual migration, running a test failover will not stop the replication or power-off the source VM. You can select to use the predefined network settings or connect the VM to a specific network. Once a test failover task is completed you can login to your VMs at the destination side and make sure your applications are up and running. The test migration can be executed at any time before the actual migration that will allow you to test your migration plan. Once you are done with the Test failover remember to run the **Test Cleanup** to remove the destination test VMs and free up resources,
- **Migration Start** – once you start the migration VCDA will perform an **Online Sync** after that it will **Shutdown** the source VM and perform one last **Offline sync** before **Power on** the VM at the destination side. The VM will be offline only while the last offline sync which should be very little due to that fact that it's done right after an online sync.
- **Migration End** – at this point VCDA will register the recovered VMs at the destination Cloud Director side and apply the predefined network and guest customization setting. Once you make sure that migration is successful you can delete the migration from VCDA.

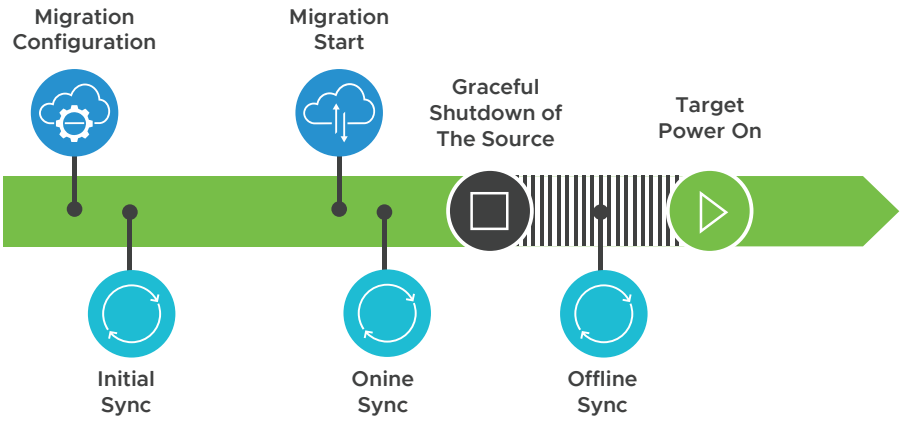

Figure 3: Migration workflow overview

L2 VPN

As mentioned earlier, a tenant may use VCDA 4.2 and above to extend a given Layer 2 network from its on-premises environment to a cloud. One may find themselves asking – What's the benefit of going through the process of configuring and deploying a L2VPN for my environment? There are many use cases which would necessitate the introduction of a L2VPN. A customer may have legacy software which requires Layer 2 network adjacency, and the VM which hosts that software needs to be moved to the cloud.

Another example may include the usage of critical applications which require a few (if any) moments of downtime. For that kind of scenario, the presence of a L2 VPN will prove to be advantageous for the ability it provides Operation Teams to stagger the migration to the cloud of VMs that make up a given critical application. Let us look at what a typical migration scenario would be for a critical application from an on-premises environment to a Cloud environment. See **Figure 4** for a high-level diagram of such a move.

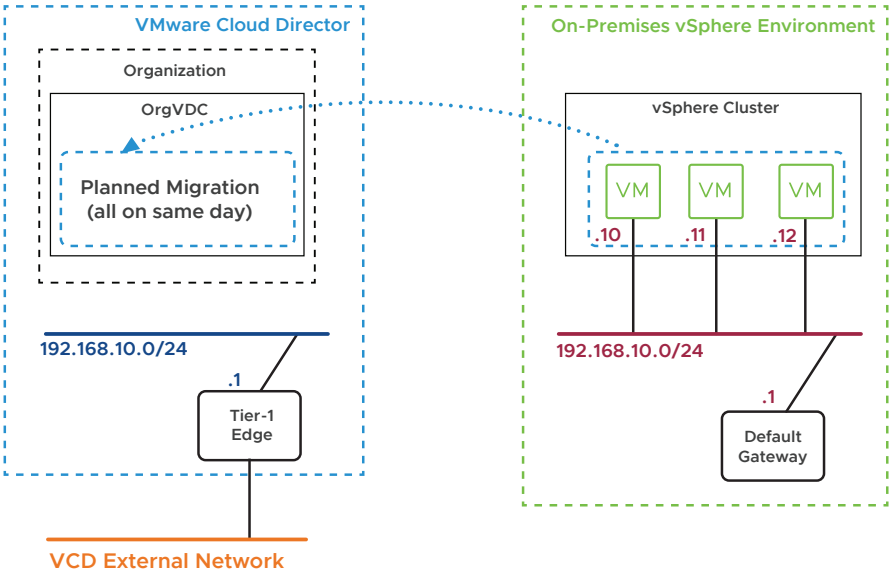

Figure 4: Migration without L2 and re-IP

- Data representing the VMs that make up the application get moved, replicated, or both to the Cloud Provider.
  - VMs are "brought up" and powered on in the Cloud Provider.
- Network connectivity from the VMs in the Cloud to critical services is tested.
- Network connectivity from the VMs in the Cloud to the on-premises environment is tested.
- Application representatives from various development groups begin validation tests to confirm overall functionality.
- Based on the outcomes of those validation tests, the application continues to stay live in the Cloud...
- ... or rollback procedures begin.

This scenario requires the presence of many personnel in a company, and the rollback process may prove to be expensive in terms of downtime of an application. Typically, with so many different people involved in such a massive move, mistakes can occur. What if there was a way to split up the work needed to migrate to the cloud, so that different parts of the application can be moved at different times? That is where the L2VPN can help. For the above workflow of the migration, VCDA 4.2 and a L2VPN tunnel can simplify things into the following list of tasks. See **Figure 5** for a high-level diagram of such a move.

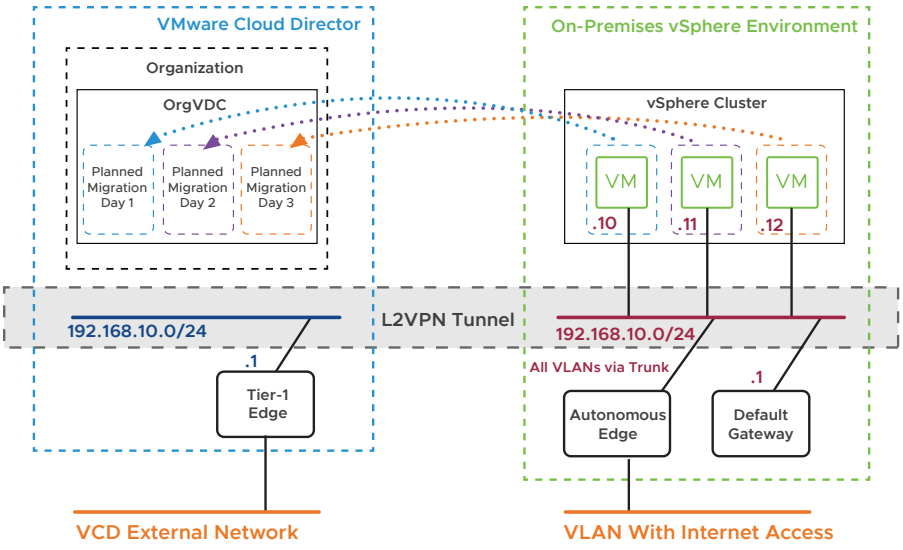

Figure 5: Migration with L2 Stretch

- Data representing the VMs gets replicated for days/weeks prior to the move to the Cloud Provider.
  - Specific Recovery Point Objectives can be set for different Virtual Machines, based on criticality.
- Migration of a specific VM is planned for the first of three weekends.
  - VCDA automatically brings up the VM in the Cloud Provider environment.
- All VMs that are not to be migrated on this first weekend will still reside in the on-premises environment.
  - Only the personnel related to the VM that was migrated need to be involved for testing.
- Because the tunnel represents a singular Layer 2 broadcast domain, network connectivity should mimic what existed prior to the migration.
- Based on the outcomes of the validation tests, the VM continues to stay live in the Cloud...
- ...or in the case of a rollback, only the VM that was migrated needs to migrate back to the on-premises environment.

Once all VMs in the application are successfully moved, then the network team can take part in the much easier task of rolling over routing for the subnet to go to the cloud. This is compared to doing that in addition to troubleshooting application connectivity across multiple operational teams and development groups.

The presence of the L2VPN enables teams to migrate to the cloud safer, and without having to do everything in one maintenance window. This reduces the failure domain of the application during the period of the migration, eases rollback activities in the case of failure, and allows only relevant personnel to be involved for potentially late-night migration activities.

Item No: VMware-Cloud-Director-Availability-Migration-WP 11/21