VMware Cloud Director Availability™ 4.1

The Natural Partnership

AT A GLANCE

VMware Cloud Director Availability offers simple, secure, and cost-effective onboarding, migration, and disaster recovery as a service (DRaaS) to or between multitenant clouds. 4.1 brings updates to the existing 4.0 platform in terms of coverage, control and breadth, to be covered in the next sections.

KEY BENEFITS

Simple Operation for DR and Migration

A converged solution with simplified recovery workflows, unified management and onboarding using familiar tools. Fully integrated with VMware Cloud Director providing intuitive and efficient DRaaS capabilities.

Simple customer driven deployment of an appliance and configuration of a tunnel from on-premise in vSphere console with self-service access migration, protection and recovery for vAPPs and VMs to and between provider VMware Clouds. Whilst providers can manage the tenant service bandwidth, compute / storage capacity and service stability through integrated workflow and public API. Providers can monitor the operational status of the service with customizable notifications and events to their operational systems as well as integrated reporting and capacity planning available in the vRealize Operations Tenant App for VMware Cloud Director Availability and natively in the VCD plugin UI.

Simple Consumption

All tenants have different workload criticality and expectations of a DRaaS Service. Cloud Providers can provide tailored DR capabilities as options to tenants to choose to protect a workload even include encryption with Cloud-to-Cloud use cases. With many complexities removed with SLA profiles, it is now far simpler to protect or migrate VMs and hence drive consumption of the service.

Tenants can consume more storage and services by storing longer term replicants for fast retrieval outside of the normal DRaaS cycle, driving more convenience and longer age of recovery. Tenants can also see their consumption of storage and other metrics to help them understand their consumption. With the Tenant App, there is more visibility into the service and if a provider has configured it, costs.

Integrated solution

The service is fully integrated with VMWare Cloud Director 10.2, now in-content protection status and jobs can be run without going into VMware Cloud Director Availability. NSX-T 3.0.1 and 3.0.2 are now integrated with VMware Cloud Director Availability, aligning with the direction that VMware Cloud is taking with NSX.

Operationally things are now easier with integrated syslog output that can send systems status and tenant issues directly to a syslog server for processing. Integrated events are also available to tenants as email, the provider can choose what is permitted for a tenant to subscribe to and tenants no longer need to be looking at the product to understand and stay informed on events and issues.

API functionality is consistently extended to cover more metrics and tasks to enhance integration and automation coverage. The Public swagger API functions include many examples, bringing more into the public specification to encourage automating DRaaS operations. Also, operations like Backup and Restore have been automated in a new user interface and available via API.

Key New Features and Capabilities in 4.1

Improved Coverage

As VMware Cloud Director continues to grow in capability, there are new versions of the core platform to support, 10.2 and NSX-T 3.0.1 and 3.0.2. Interoperability with these versions of products is now covered in VMware Cloud Director Availability 4.1.

As partners grow their customer base of DRaaS consumers, it has become necessary to be able to provide filtering for situations where there are 100's of customer orgs on a provider Virtual Datacenter (pVDC). Also, in consideration of growth, in 4.1 where VMware Cloud Director is managing vCenter replication in multiple geographies, the provider can now place the VMware Cloud Director Availability components nearer to the vCenters to reduce the number of traffic paths to manage, whilst still managing from a single instance.

The other area of coverage improvement is the API and the number of capabilities that the public API supports has in 4.1 expanded to include a backup and restore function, initial configuration and peering configurations to assist with management and operations. Lastly, examples are now included of common use cases to help jumpstart development.

Improved Controls

In 4.0 we provided provider only notifications via Syslog, now in 4.1 notifications are available to tenants and providers by email as well. Tenants have the ability to configure their notifications and events they wish to see via email, or a provider can lock down via policy control what the tenant can configure. This is particularly useful if providers need to filter events that may cause more support than is necessary.

For Providers, syslog to the VMware Cloud Director Availability content pack for vRealize Log Insight gives them the essential event and escalation capability to manage the service, this is now complimented with configurable and automated email notifications should they need to inform non-operational personal of events.

Configuration controls have been enhanced with backup and initial configuration:

- Backup operations now automate the backup of replications and their status and all services; Management Service, Tunnel Service and Replicator Services to password protected archives. This can be achieved via API or via the new backup menu in the UI.

- Initial configuration of VMware Cloud Director Availability replicator and tunnels used to be achieved with a separate UI for each, now this is all achieved in a single UI wizard directly in the management appliance - adding Replicator Service instances, Tunnel Service details, lookup service details used by all appliances and future proofing for certificate management to come in later releases.

Improved Offering scope

Keeping customer's security cloud requirements in mind, we now support the use of encryption polices in Cloud-to-Cloud replication and encrypt the replications at the target site, providing the source and target are using the same KMS server cluster. This has limited interop at this stage to VMware Cloud Director 10.1.1 and 10.2 and VC 6.7U3 and limited capabilities in disk flexibility.

Multiple VMware Cloud Director authentication mechanisms are now supported, and customers can use their own identity provider SSO / LDAP / SAML, etc. in order to instantly login across the associated orgs in a multi-site VCD environment. In the Cloud-to-Cloud use case this makes for a quicker configuration as you do not have to log in to each organization every time you create a replication.

Lastly in terms of offering, providers are now able to enable or disable Migration and Disaster Recovery replication separately for further granularity to service tiers for the customer offering.

Disaster Recovery Service Organization Controls

- Whitelisted, activation-controlled service is disabled by default, allowing partner to upsell DR to customers. Org controls include max RPO, snapshots, and replicants to help tier tenant DR offerings.
- VMware Cloud Director Availability is integrated into VMware Cloud Director to provide administrators at a glance protection status and in-context menu to start and manage replication and DR actions to protect vApps and virtual machines.
- Service Providers can now define and control Service Level Agreement (SLA) replications settings: Migration and/or Protection, Recovery Point Objective (RPO), retention policy for the point-in-time instances, quiescing, compression, and initial synchronization time by using SLA profiles that can be assigned to organizations making it simpler for an organization to consume the DRaaS.
- To restore a workload to a previous state, tenants can use point-in-time or stored instances. To avoid the automatic retention of point-in-time instances, tenants can store point-in-time instances (the number is limited by the provider in the SLA profile for the org). The stored instances do not change, you can use them to recover the workload to the stored instance, regardless of the overall retention period of the point-in-time instances.
- Service providers can configure a limit for the replication data traffic from the on premises to cloud sites for regions where bandwidth is a premium. Providers can also set a global limit in VMware Cloud Director Availability for the total incoming replication traffic from all cloud sites.
- For larger environments provider datacenters can have their own VMware Cloud Director Availability instance mapped and working with the Provider Virtual Data Center(s) in the data center. Multi-geo tenants can have separate connections to each regional datacenter from their central site if required. Also, multiple geo-sites can which are managed by a single VCD (site associations supported in VCD 9 onwards) are now supported by VMware Cloud Director Availability where each instance in each geo can be logged in without re-authenticating. This capability allows customers to use their own identity provider (SSO / LDAP / SAML etc) which is critical for certain verticals and security sensitive customers who do not want to use a DR local user account.

On-Premises to Cloud and Cloud-to-Cloud Disaster Recovery

Tenants can completely self-serve and automate replication, migration, failover, and failback of VM and vApps and post failover operations from their vCenter plugin or from the VMWare Cloud Director interface using the solutions symmetric capabilities. The user interface provides better usability and efficiency of easier replication management and overview of the tasks by simplifying the management interface.

- Service providers must be able to resource plan to ensure that the service can satisfy any failover testing or events within their clouds. Resource requirements for compute vCPU and vGB RAM can be listed and aggregated for customer org VDC, organizations and provider VDC to help with capacity planning. This is very important for providers to manage as customers are not limited to the number of tests that they can execute.
- Storage management is key for Disaster Recovery as customers replicate their workloads and variations in RPO and the number and granularity of multiple point in time instances will dramatically impact the amount of storage used. VMware Cloud Director Availability now shows the disk space used by each virtual machine replication and aggregates the disk usage information per organization. Tenants can monitor the disk usage for every replication in all directions to help them understand each workloads' requirement and providers can monitor the disk usage for every organization via UI and API for metering and capacity management.
- Storage encryption policies that are available in VCD 10.1 and onwards are used to encrypt VMs using existing vSphere storage policies encryption flags. For Disaster Recovery between 2 Org VDC (Cloud to Cloud DR), encryption of the VM and the replicants (cycled multiple point in time instances) are also encrypted providing the same KMS server is used for source and destination clouds.
- What if replications fail or there are issues with the system? How is operations informed? Service Providers can use the syslog feature to monitor the event notifications that VMware Cloud Director Availability generates either by using a syslog server or in VMware Cloud Director by monitoring the VMware Cloud Director Availability events. Syslog events can be forwarded to a log manager, like vRealize Log Insight (which is included in the Flex core bundle) where they can be filtered and changed into alarms and notifications.

- What if replications fail or there are issues with the system? How is the tenant informed? Tenants can be informed of status changes in their service using email notifications, which they can customize, and the list of options is controllable from the cloud provider by policy. The provider can control what events the tenant can escalate to email notification and most likely will want only service affecting events to be sent, rather than other events that could be just 'noise' around activities happening in the service.

Migration On-Premises to Cloud and Cloud-to-Cloud

Migration to Cloud is literally the process of moving resources such as virtual machines (VMs) or virtual applications (vAPPs) from one place (an on-premises site or a cloud site) to a cloud computing environment in this case running VMware Cloud Director Availability. Migration of a VM can be accomplished in 2 ways, cold or warm migration, both have the option to use network connectivity at layer 2 which de-risks migrations and can help extend the duration of a migration.

Cold Migration

Cold migration is the migration of powered off or suspended virtual machines between hosts across clusters, data centers, and vCenter Server instances. This is why this approach is considered to be suitable for **non-business critical workloads**. The simplicity of doing such migrations makes them ideal for self-service, where the tenants execute all the operations by themselves in the vSphere console with the DRaaS plugin.

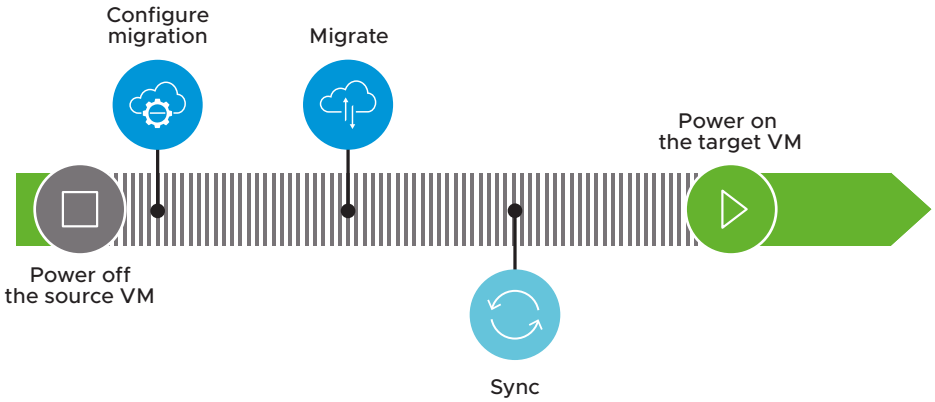

Warm Migration

With a warm migration the VM doesn't have to be powered off at migration time. This reduces downtime significantly and is one of the main reasons it is ideal for **mission-critical workloads**. However, the migrated VM typically has new network settings that might require additional actions on the target side once the process is complete, this can be negated with a layer 2 stretch (see next section). This method is still suitable for self-service, even though providers who have access to their tenants' on-premises environments can offer it as a managed service.

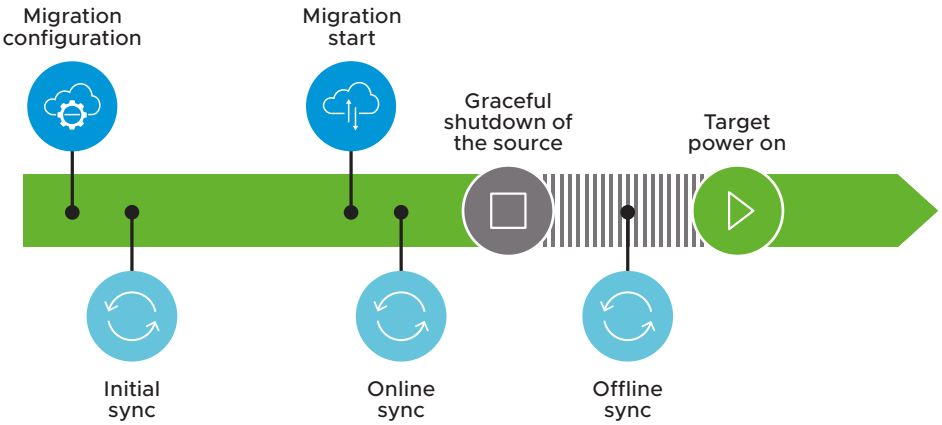

A warm migration consists of the following steps:

1. Configure Migration

   Once the migration is configured in VMware Cloud Director Availability through the DRaaS vCenter plugin, or directly, in the VMware Cloud Director Availability UI, the initial data sync will be complete. It has no impact on the source workloads, and they continue running without any interruptions. When the initial sync is complete, the workloads are ready to be switched to the destination site. During this 'wait period', changes are synchronized once every 24h.

2. Start Migration.

   After the preparation for the migration is done, it can be started by the user. At the time of switchover, online sync is executed first, followed by a graceful shutdown. If the graceful shutdown fails with a timeout, a forced power off will be triggered. When the machine is offline, a rapid sync is performed to capture any changes since the previous one.

3. Power on the migrated VM

   The final step is to compose the VM and power it on at the destination site. This sequence of actions minimizes the VM/vApp downtime to almost that of a graceful restart

Warm migrations with L2 stretch

The warm migration with Layer 2 network stretching is a little more demanding in terms of requirements, but still the most commonly used process for migrations. Extending the on-premises network to Cloud is a popular approach, and in this case, it removes the necessity to reconfigure the VM, resulting in a shorter downtime period. It also provides the capability for migrated VMs to still reach back to the source site, meaning not all components of an application need to be migrated as one job, time can be taken and the need for accurate assessment is less.

In terms of flow, it is the same as a warm migration but with an extra prerequisite – to set up Layer 2 network stretch. This method is suitable for providers to offer as a service to their tenants due to the additional requirements involved in this configuration, it is not out of the box like native migration but is relatively simple to configure.

Guaranteed Service Level Agreements

Speed of recovery and granularity of recovery are important DR and migration considerations. However not all workloads require a fast Recovery Point Objective (RPO) as the faster the RPO the more storage will be consumed. Most applications will be fine with a 1hr RPO, and VMware Cloud Director Availability will support a max. RPO of 5 minutes (in vSphere 6.5 and above), but the RPO is completely customizable by the service provider for your tenants in SLA profiles. The Recovery Time Objective (RTO) is still the time it takes for the VM to power on.

Deep Integration with VMware Stack

- Core component of Cloud Provider Platform enables cloud providers to offer differentiated services.
- Support is available for cluster datastore allowing you to perform storage migration to a cluster datastore.
- Support for edge clusters in Cloud Director ensures optimum performance of Cloud Director environments.

Service Provider "Day 2" Operations and Monitoring

- Policy based management of the DR service provided to customers is controlled by SLA profiles making it far easier to operationalize and to meter consumption and usage.
- Migrate tenants from one Cloud Director to another, for example, to set up a new data center or if there's a need to perform maintenance.
- Monitor service performance and availability using the syslog configuration and sending events to vRealize Log Insight, a default component in the 7-point flex core.
- Monitor tenant usage of storage and compute to understand your ability to recover tenants' workloads, native in the UI or over the API.

Reporting usage

Automatic metering for monthly usage is supported using 3.6.1 or higher of VMware Usage Meter, capturing replications for DR and for migration (although VMware capture migration replications, they are not charged). Service providers should deploy and configure VMware usage meter 3.6.1 H3 and above or 4.2 to take care of VMware's charges, providers can now use the API or CLI for their own metering if required as well as the Tenant App plugin for vRealize Operations.

Upgrading

VMware recommends upgrading to version 4.1 as a direct upgrade from 4.0 and the components should be upgraded in sequence starting with the

Cloud Replication Manager appliance, then the Replicator appliance(s) and lastly the tunnel appliance. (please see product documentation for instructions)

Interoperability

VMware Cloud Director Availability is compatible with:

- vSphere 6.5 U3, 6.7 U1, U2 and U3 and 7.0
- VMware Cloud Director 9.5, 9.7, 10.0, 10.1, 10.1.1, 10.2
- NSX-V 6.4.6
- NSX-T 2.5.0, 2.5.1, 3.0.0, 3.0.1, 3.0.2
- vCloud Usage Meter 3.6.1, 4.2, 4.3

Support

VMware makes Subscription Services Support available to all VMware Cloud Director Availability customers. Support includes access to specialists who assist in coordinating onboarding activities as well as ongoing service support. For customers who require additional services, VMware also offers professional services engagements on best practices and getting started with your deployment, both directly and through an extensive network of certified professionals.

How is it Sold?

VMware Cloud Director Availability is offered as a subscription-based consumption model. Units are protected VMs/month, migrations, although metered, are at zero cost.

Please note that VMware vSphere Replication is included, without additional cost, in all supported VMware Cloud Provider bundles that contain VMware Cloud Director. For service provider system requirements and interoperability, see Cloud Director Availability documentation.

LEARN MORE

To Learn more about how VMware Cloud Director Availability 4.0 works, please visit cloudsolutions.vmware.com or please watch and subscribe to our YouTube Channel or any of the resources below:

http://bit.ly/VCPP_Twitter

http://bit.ly/VCPP_FB

http://bit.ly/VCPP_LinkedIn

http://bit.ly/VCPP_Youtube

FOR MORE INFORMATION OR TO PURCHASE VMWARE PRODUCTS

call 877-4-VMWARE (outside North America, +1-650-427-5000), visit *http://www.vmware.com/products*, or search online for an authorized reseller. For detailed product specifications and system requirements, refer to the documentation.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.

Copyright © 2019 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: vmw-ds-temp-word 2/19