

VMware vCloud® Availability 3.0

Solutions Overview

Q. What is VMware vCloud® Availability 3.0?

A. VMware vCloud Availability is a powerful solution used by VMware Cloud Providers to offer simple, secure and cost-effective onboarding, migration and disaster recovery services “to” or “between” multi-tenant VMware clouds.

Q. What capabilities does the solution provide?

- Self-service protection and failover workflows per virtual machine (VM) and vApp
- Single installation appliance installation simplicity
- The capability of each deployment to serve as both source and recovery vCloud Director instance (site). There are no dedicated source and destination sites.
- Symmetrical replication flow that can be started from either the source or the recovery vCloud Director site means the UI can be accessed from anywhere with correct context
- Replication and recovery of vApps and VMs between vCloud Director sites
- Maximum replicated VMs and retained snapshots as well as minimum RPO Policy controls for providers to apply to one-to-many VDC. This helps to control storage costs and provide tiered services to customers
- Secure Tunneling through TCP proxy between sites
- Integration with existing vSphere environments
- Multi-tenant support native in vCloud Director
- Built-in encryption and compression of replication traffic
- Build your own failover plan by JSON import or use existing out-of-the box vRealize Orchestrator workflows.
- Support for vCenter Server versions 6.5 U1 and 6.0 U3, and vCloud Director versions 8.20.X, 9.0.X, and 9.1,9.5

Q. What use cases does VMware vCloud Availability 3.0 support?

A. On-prem to Cloud migration, On-prem to Cloud DR, Cloud-to-Cloud DR, Cross version vCD migration.

Pricing and Packaging

Q. For Cloud providers: How is VMware vCloud Availability 3.0 packaged and how may it be purchased?

A. vCloud Availability 3.0 is available as a Pay-As-You-Go service to Service Providers in the VMware Cloud Provider Program. The service is metered on a monthly basis based on number of VMs protected. Detailed information is available in the VMware Cloud Provider Program Guide at <http://www.vmware.com/partners/service-provider.html>

Q. I am a Cloud Provider and am not currently enrolled in the VMware Cloud Provider Program. Can I purchase this product directly from VMware?

A: You must be enrolled in the VMware Cloud Provider Program in order to purchase vCloud Availability 3.0. To learn more about the VMware Cloud Provider Program, please visit <http://www.vmware.com/partners/serviceprovider.html>.

Q. For Enterprise Customers: How is VMware vCloud Availability 3.0 packaged and how may it be purchased?

A: Enterprise customers must consume this offering through a VMware Cloud Provider Partner that is offering this service in their vCloud Director VMware clouds. All prices for these services will be quoted by the VMware Cloud Provider Partner. To find a VMware Cloud Provider Partner offering DRaaS please use the assisted search here: <https://cloud.vmware.com/providers/guided-search>.

Migration features

Q. What is the VM/vAPP migration feature?

A. vCloud Availability allows end users to replicate and migrate virtual machines and vApps from on-prem to VMware Cloud and between different VMware Cloud environments. End users can set-up an organization virtual data-center on destination and migrate virtual machines from source data-center in few simple steps including assigning destination networks. This provides for a predictable way to migrate workloads in a self-serving manner.

Q. How is failover (migration) performed?

A. To perform a migration the virtual machine or virtual app must be protected (replicated) between the source and target location. Once replicated, an optional resync can be initiated prior to failover (migration) to get real-time data migrated before failover to the destination site. This ensures that the latest changes of the source vApp/VM are present in the recovered instance.

- If you are migrating a vAPP, it is important to manage the state of all VMs in the vAPP for a stable state; this can be validated simply in the UI and corrective actions can be taken.

Q. How are VMs initially replicated between clouds?

A. Once the workflow is configured to start VM or vAPP replication the VMware vCloud Availability vApp Replication Manager ensures that delta information is sent from one ESXi host to another ESXi host. Management and Monitoring information for the replication is available from the vApp Replication Manager portal and APIs.

Q. What is a test failover?

- A. Test failovers allow you to verify whether the source data is replicated correctly on the destination.
- You can test network connectivity and application (VM) behavior
 - Seed data is supported, as is sync, prior test failover
 - vAPP can also be powered on to test

Q. What actions can be taken by users?

- A. By using the Actions pane in the DR Workloads page, you can perform the following tasks:
- Failover workloads among to-destination sites
 - Failback workloads among from-destination sites
 - Failover Reverse workloads to synchronize data between source and destination sites
 - Failback Reverse workloads to synchronize data between source and destination sites
 - Test replication tasks and Cleanup test data

Q. What functionality is available to monitor DR operations?

- A. You can monitor the overall vCloud Availability status by using the vCloud Availability Portal home page in vCloud Director or in the separate UI and view the progress and status of finished and ongoing tasks by using the Replication Tasks tab.

Policy control features

Q. What policies can be used to control / limit functionality?

- A. A Cloud Provider can assign replication policies to local one-to-many VDC organizations:
- Policy to assign an vCD organization as a replication source and/or destination
 - Minimum Recovery Point Objective (RPO)
 - Maximum number of VM replications
 - Maximum number of point-in-time instances per VM replication

Following are the default policies available in vCloud Availability. Users have the option to modify these default policies:

- Allow outgoing replications = false
- Allow incoming replications = false
- Min RPO = 15 min
- Max instances = 24
- Max replications = 0
- The source org allows outgoing replications
- The destination org allows incoming replications
- The replication RPO
- The number of instances
- The total number of destination replications

Architectures

Q. What services are included in the installation; do I need to configure vSphere Replication somehow?

- A. This is where a lot of improvements have been made. The architecture of the solution uses a vCloud Availability Replicator, a vCloud Availability Replication Manager and a vCloud Availability vApp Replication Service/Manager together to support replication, secure communication and storage of the replicated data.

Each service provider can support recovery for multiple customer Org Virtual Data Center environments that can scale to handle increasing loads for each tenant, and for multiple tenants.

All replication matters are handled within the vCloud Availability user interface (or API) and are simple workflow configuration driven tasks for one-to-many VMs and vAPP.

Minimum requirements

Q. What are the minimum requirements for vCloud Availability 3.0?

- A. For system requirements and interoperability, see vCloud Availability documentation.

Service Deployment

Q. Does vCloud Availability 3.0 require any agents to be deployed?

- A. No, the solution is agentless and uses vSphere Replication.

Q. How is the product installed in the provider data center?

- A. vCloud Availability can be deployed using the VMware OVF Tool. Alternatively, you can use the vSphere Web Client to install the vCloud Availability service; all DR services are deployed via a single installation OVA package:
- Combined install – all in one deployment option for test / eval environments
 - Manager node with vCloud Director support – set configuration for compute and storage
 - Varying sizes of replicator node based on your capacity needs:
 - Replicator node - set configuration appliance of 2 vCPUs, 4 GB RAM, and 10 GB storage
 - Large replicator node – set configuration appliance of 4 vCPUs, 6 GB RAM, and 10 GB storage
 - Tunnel node – set configuration appliance of 2 vCPUs, 4 GB RAM, and 10 GB storage

Q. How is the product installed in the on-prem data center?

- A. The on-prem appliance is an OVA that can be installed via UI or OVF Tool. It is a single appliance that has all the required components embedded to facilitate the connection and replication with Cloud Providers.

Q. Why are there different versions of replicators?

A. Replication in terms of volume will impact the capacity and performance of the appliance. When each VM is compressed and encrypted there is an overhead on CPU. Whilst encryption is mandatory, compression can be optional, and both tax system resources.

Multiple replicators can be added to your DR environment to suite processing needs.

Q. What are the tested scale limits for a deployment?

- 4000+ replications across 100 tenants
- 300 active protections using a single large vCloud Availability Replicator. For more information about the replicator types, see Deploy vCloud Availability for Cloud-to-Cloud DR Services by Using the vSphere Web Client.

- 100 tenants with active replications
- Up to 4 TB VMDK limit (tested)

Resources

Q. Where can I find more about vCloud Availability 3.0?

- For more information visit: <https://www.vmware.com/products/vcloud-availability.html>