

Architecting VMware Cloud Director Availability Solution

A Natural Partnership

For Cloud and Service Providers

Atanas Stankov

Senior Solutions Architect / VCPP Engineering

February 2021



Table of Contents

Introduction..... 4

Use Cases 4

Disaster Recovery..... 4

Migration 5

VMware Cloud Director Availability Architecture 5

VMware Cloud Director Availability Components 5

Network flow7

VMware Cloud Director Availability Management..... 8

Cloud service..... 8

Manager service 9

VCDA Portal..... 9

VCDA API 11

VCDA extension in VCD Portal 11

VMware Cloud Availability Data Path 13

Tunnel..... 13

Replicator 13

MultiNIC appliances 14

Relations of other infrastructure components 15

Platform Services Controller 15

vCenter Server 16

VMware Cloud Director 16

ESXi hosts 16

Initial service configurations..... 16

VCDA 4.1 and newer17

VCDA 4.0.x and earlier17

Replicator17

Replicator Manager17

Tunnel.....17

vApp Replication Manager 18

Public API Endpoint..... 18

Registering VCDA extension in VMware Cloud Director 18

Enable Tunneling 19

SSO Registration 19

Pairing 19

- General information 19
- Pairing over Internet 20
- Pairing over private networks 20
- Admin access from Internet 20
- Throttling 20
- Policies 21
- SLA profiles 21
- VMware Cloud Director Availability at on-premises site 22
 - Overview 22
 - Requirements 23
 - Configuration 23
 - Pairing 23
 - Placement 24
- Appliance management 24
 - Network configuration 24
 - NTP 24
 - Certificate Management 24
- Monitoring and reporting 25
 - Built-in reporting 25
 - UM integration 26
 - VRLI pack 27
 - VROPs pack 27
- List of Figures 28
- Reference Documents 28
- Acknowledgement 29
 - Author 29
 - Reviewers 29

Introduction

VMware Cloud Director Availability is a Disaster-Recovery-as-a-Service solution. It provides protection and migration for VMware Cloud Director vApps and virtual machines. VMware Cloud Director Availability, aka VCDA, is available for participants in VMware Cloud Provider Program and allows them to protect and migrate vApps and VMs:

- From on-premises vCenter Server site to a VMware Cloud Director cloud
- From VMware Cloud Director cloud to an on-premises vCenter Server environment
- Between VMware Cloud Director managed clouds

The goal of this document is to provide information how to architect a DRaaS solution for VMware Cloud Director based on VCDA 4.0.

Note: VMware Cloud Director Availability and vSphere Replication are two different and independent products and there is no interoperability between them. Both products can coexist in a single vSphere infrastructure but it's not possible to replicate and migrate virtual machines between them.

Use Cases

VCDA supports two different use cases – disaster recovery and migration of vApps/VMs. Both of them rely on replication of virtual machines. In both cases at least one of the sides is a VMware Cloud Director, aka VCD, managed cloud and the other side could be another VCD cloud or vCenter Server. VCDA cannot protect or migrate bare metal servers or VMs managed by non-VMware hypervisors.

Disaster Recovery

A tenant can purchase DRaaS provided by a public cloud and based on VCDA to protect its virtual machines running in its on-premises datacenter. Cloud provider will assign a portion of compute, storage and network resources from its own cloud and will group them in an Organization Virtual Data Center (OrgVDC) to the tenant. Also cloud provider will enable this OrgVDC (respectively the tenant) to protect its on-premises virtual workloads to the cloud with specific parameters like minimum RPO, bandwidth control, direction of protections and other. This will allow tenant to failover its VMs in case of outage in its on-premises datacenter. With VCDA 4.0 tenant can configure a number of settings for a VM before its initial synchronization – what RPO will be (minimum is 5min), if multiple point-in-time copies will be preserved for certain amount of time, if the same VM has a cold copy at destination will it be used as seed VM to save time and network bandwidth during initial sync and few other configurable settings.

As some tenant may decide there is a lot of complexity and may ask to have a simplified workflow cloud provider may use so called SLA profiles. SLA profiles are used to control majority of these settings and predefine them for a specific tenant. In this way settings from SLA profile are pre-populated and tenant does not need to configure them for every single protection.

There are also settings that control how VM will be configured after it is failed over – to which network will be connected each network adapter, choose how to configure IP address of each network interface, if there are multiple point-in-time (MPIT) copies will they be consolidated during failover operation or manually later, which MPIT will be used during failover and so on.

If cloud provider offers managed services and tenant accepts such offer it's possible for cloud provider to take care and configure all protections and after that to transfer ownership to tenant.

After a replication is configured and full sync - initial or manual, is completed there are different operations that can be performed over each protection:

- Manual sync – administrator can initiate manual sync out of configured PRO replication window
- Test Failover – used to test if a vApp/VM could be successfully created from replicated data and powered on in destination side. Original vApp/VM at the source side stays operational. Usually, this operation is followed by Cleanup operation which removes vApps/VM created during the test at destination. Replications continue during the tests instances at destination exist according to RPO settings
- Failover – used when source side is down. At destination side a vApp/VM is constructed from replicated data and is powered on.
- Reverse – after a successful failover to destination side and when the source side is again up and running tenant may decide to reverse the replication direction from original destination where VM is up to data to the original source site.

This will provide a DR protection for this VM again. It's possible to use original VM from source side as a seeding VM and this will move between sides only changes blocks.

- Migrate – both sides are up and running. When this operation is initiated vApp/VM at source side is powered off, in destination new vApp/VM is constructed from replicated data and is powered on.

VCDA allows to perform operations not only over a single vApp/VM but also over a bunch of vApps/VMs. For more details what concurrency is supported for each operation please check Release Notes for the version you run.

Migration

When tenant plans to do planned migration of workloads to the cloud it may use Migration workflow which simplifies the process. When a “New Migration” is configured VCDA starts a replication of vApp/VM from source to destination. When initial sync is completed source VM is powered off and any delta generated in the source site is synchronized to the destination. At destination side VCDA requests from VCD to construct a new vApp/VM from replicated data and powers it on. Migrated workloads could be customized to make them ready for operation in the new virtual datacenter.

As with Disaster Recovery use case, Migration can be configured for groups of multiple vApps/VMs.

Each VCDA operations generates a number of operations in VCD and vSphere infrastructure layers. At both layers different objects are created, intensive read and write operations are triggered, network bandwidth utilization increases. Before a batch of concurrent operations is initiated it is required to understand how this will impact performance of computing resources, storage devices and network connections. Bulk operations should be executed with understanding how the other workloads will be affected.

To address this possible negative effect VCDA provides policies which help cloud providers to control concurrency and respectively the impact over the cloud infrastructure.

Migration use case could be also offered by cloud provider as a managed service.

VMware Cloud Director Availability Architecture

In the cloud, VCDA requires existing VMware Cloud Director installation for be fully operational. This means VCD has to be installed and configured and ready to provide tenants with compute resources through Organizational Virtual Datacenters.

VCDA site concept in the cloud consists of the following set of appliances:

- One VCDA Manager appliance
- One VCDA Tunnel appliance
- One or more VCDA Replicator appliances

For tests purposes it's possible to deploy a combined appliance in the cloud where all services are hosted on a single appliance.

Note: Combined appliance is not supported for production environments. It's supported only for labs and small Proof-of-Concept (POC) installations.

In the on-premises datacenter a single appliance is deployed and configured to replicate or migrate virtual machines running in vCenter managed infrastructure.

VMware Cloud Director Availability Components

The diagram below shows components that build VMware Cloud Director Availability solution.

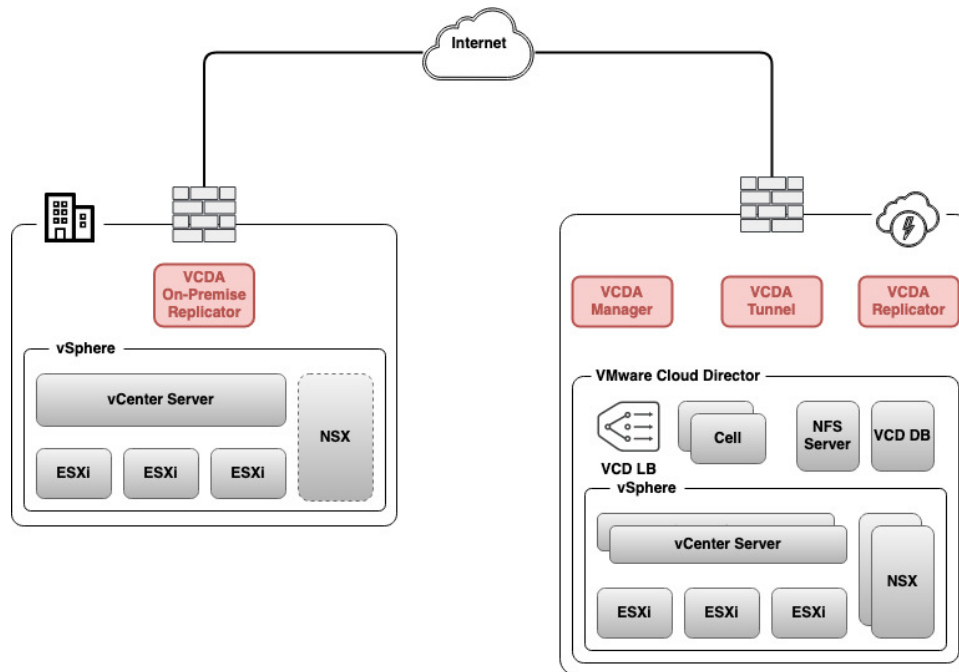


Figure 1 - VMware Cloud Director Availability Components

The grey components are pre-existing before VCDA is introduced in the cloud or in the on-premises datacenter. The red components are VCDA components. The diagram shows the bare minimum of components required to build fully functional disaster recovery solution based on VCDA.

The table below briefly describes the role of each VCDA component:

Component	Role
Tunnel	<p>Tunnel appliance is the single-entry point to VCDA instance in the cloud and its role is to handle incoming management and replication traffic. Tunnel handles both data and management traffic and forwards it respectively to cloud replicators and manager. No other VCDA component needs to have a dedicate Internet accessible endpoint.</p> <p>Tunnel is a mandatory component since version 3.0.</p>
Manager	<p>Manager is responsible for communication with VMware Cloud Director and through this communication VCDA discovers resources (OrgVCD, storage policies, datastores, networks, etc) managed by VCD and used by tenants. This information is required to discover vApps/VMs that could be replicated/migrated or suitable locations for destination of incoming replications/migrations.</p> <p>It also provides UI and API interfaces to use VCDA and interact with it. Another role of the manager is to communicate with local and remote replicators and receive data from them about each protected/migrated workload.</p> <p>These functionalities are realized by two VCDA services that run in Manager appliances – cloud.service and manager.service. Both services are explained in next chapters of this document.</p>
Cloud Replicator	<p>Cloud Replicator is responsible to move replication data to and from ESXi hosts in the cloud.</p>

Component	Role
	<p>For outgoing replications/migrations it communicates with VMKernel interface of an ESXi host and captures and encrypts replication data, optionally compresses and encrypts this data and sends it to remote replicator which can be another cloud replicator or on-premises replicator.</p> <p>For incoming replications/migrations cloud replicator receives replication data from a replicator (cloud or on-premises), decompresses and decrypts this data and sends it to ESXi to be written on a datastore.</p> <p>Cloud Replicator is the only component that can scale out as the number of protections/migrations increases.</p>
On-premises Replicator	<p>On-premises replicator is deployed in tenant on-premises datacenter. It creates a pairing relation to VCDA in the cloud and can protect and/or migrate VMs running locally to the cloud and vice versa.</p> <p>On-premises replicator does not require public endpoint as it only initiates connectivity to the cloud.</p> <p>A single on-premises replicator can protect VMs from a single SSO domain in on-premises even if there are multiple vCenter Servers in this SSO domain. If the requirement is to do replications from the cloud to on-premises, then a single replicator is required for each vCenter due to the way how placement works.</p> <p>Also, one on-premises replicator can be paired with a single VCDA instance in the cloud. If tenant intends to use more than one DR service in the cloud a dedicated on-premises replicator is required for each cloud DR service.</p>

Table 1 - Appliance Roles

VCDA deploys software packages only for its integration with vSphere Web Client and VMware Cloud Director Portal. VCDA does not deploy any software packages neither in ESXi hypervisor nor in VM guest operating systems. ESXi hosts are out-of-the-box ready to interoperate with VCDA to protect and recover vApps/VMs.

There is a single exclusion of this rule and it's related to the case when replications of encrypted VMs between clouds will be configured. In this case a special VIB has to be deployed on each source or destination ESXi host. More details are available at <https://docs.vmware.com/en/VMware-Cloud-Director-Availability/4.1/VMware-Cloud-Director-Availability-User-Guide/GUID-25DB8847-67ED-410D-B438-0DA470B977E5.html>.

Network flow

The following diagram shows network flows between VCDA components and the other infrastructure components on which it depends.

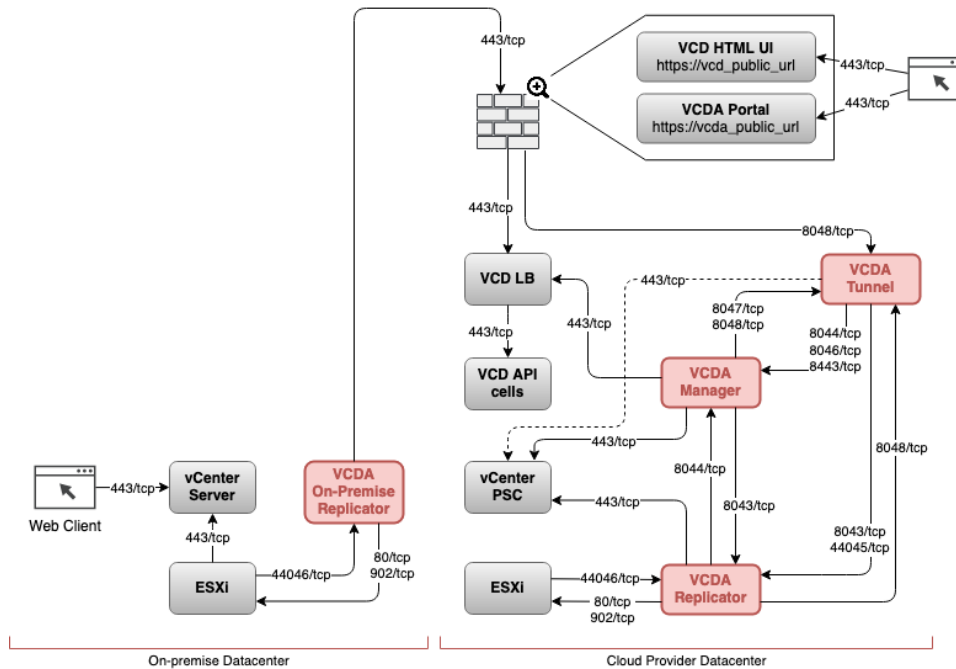


Figure 2 - Network flow in VMware Cloud Director Availability

On this diagram is shown scenario when on-premises tenant is using DR service in the cloud over Internet. It's possible also for a tenant which has a connectivity to cloud provider over private networks (VPN for example) to connect to Cloud Provider DR service over this private network and avoid sending replication traffic to Internet where links could be slower and much more expensive. This is fully supported by VCD. In such cases VCD On-premises replicator will have to communicate directly with VCD tunnel in the cloud on port 8048/tcp. Pairing over private networks between tunnels on 8048/tcp for cloud-to-cloud scenario is also supported.

VMware Cloud Director Availability Management

Management of VCD solution is completely concentrated into VCD Manager. All appliances – tunnel, manager and replicator, have appliance management interfaces, but solution is configured and managed by VCD Manager. Appliance management and its importance is discussed later in Appliance management.

Manager has low resource requirements and can be deployed in management cluster without significant impact on overall resource utilization.

From connectivity standpoint it needs to be able to communicate with:

- Lookup Service of vCenters that provide resources to VMware Cloud Director and
- HTTPS interface of VCD cells

VCD Manager is composed by two services – cloud and manager.

Cloud service

This is the service which understands VCD constructs – OrgVCD, vApps, networks managed by VCD, storage policies, etc. To achieve this VCD Manager communicates with VCD API through VCD LB. VCD Manager does not communicate with VCD consoleproxy cells/interfaces. Using this communication, VCD is able to:

- discover VCD managed vApps/VMs and protect/migrate them to another DR-enabled cloud or on-premises vCenter
- discover suitable destination for incoming replications/migrations.

Cloud service manages pairings with other DR-enabled clouds, policies and their assignment to VCD organizations and SLA profiles. Cloud service provides information for replication tasks and system tasks. Cloud service also reports what compute resources are consumed by replications per tenant and per PVDC, what the storage consumption per datastore is, plus many other high- and low-level details.

Cloud service management interface is available on https://vcda_manager_fqdn:443/admin. It's possible to login with local OS root account, SSO account if appliance has registration in SSO domain or with VCD System Administrator account if initial configuration is completed.

Manager service

Manager service is the one that has registrations of local and remote replicators. Remote replicators are registered in manager during pairing process. For each replication Replicator Manager chooses one replicator from source site and one from destination site. Destination replicator is responsible to discover suitable resources to create replica disks at destination and write data. In case of disaster manager and destination replicator have the responsibility to failover protected VM. Replicators send information to Replicator Manager about their operation – statuses, amount of data replicated, operations start time, time to complete and other.

Manager service is also used to manage replicators. It may put a replicator in maintenance mode which assign another replicator for each replication. Or it may trigger rebalance of replications across all replicators. This is useful when existing replicators need to be offloaded from some replications by adding a new replicator in the solution.

Manager service management interface is available on https://vcda_manager_fqdn:8441. It's possible to login with local OS root account and with SSO account if appliance is registered in Lookup Service.

VCDA solution management interfaces

There are several options to manage VCDA solution in the cloud and these are:

- VCDA portal
- VCDA API
- VCDA extension in VCD portal

For managing VCDA in on-premises datacenter options are explained in *VMware Cloud Director Availability at on-premises site*.

VCDA Portal

VCDA portal is served by VCDA Manager and more specific by the cloud service. In order to use the portal cloud provider is required to register a DNS A resource record which is resolved to a public IP address. Then for this public IP address cloud administrator has to configure a publishing rule to redirect traffic destined to VCDA portal to tunnel appliance on port 8048/tcp. After that tunnel forwards the request to its final destination vApp Replication Manager over 8443/tcp which responds to the request.

Internally VCDA portal is open from https://manager_fqdn_or_ip:443.

By default, Portal will load /ui/login page. It allows to login with credentials from VCD in form of <username>@<Organization_name>. Organization_Name could be also System. It's possible for administrator to use /ui/admin page where it can login with local OS root account or with account from SSO domain. Admin portal will be loaded for:

- Logins to /ui/admin
- Login to /ui/login with VCD System Administrator account

Login to /ui/login with OrgName account will load user space which has no permissions to do configuration changes and UI is simplified.

If tenant organization is configured to use SSO authentication users are not able to authenticate to VCDA Portal and they have to use VCDA UI plugin in VMware Cloud Director.

VCDA features are implemented via APIs. Some of the API calls are public and majority of the API calls are private. VCDA portal has all features implemented no matter if public or private API call is doing the job in the background.

Portal is the main management and monitoring interface of VCDA. Management functionality allows to configure appliance settings and cloud service settings. Monitoring functionalities allow to get different levels of information for VCDA solution.

Dashboard provides high-level information about the status and usage of VCDA. Screenshot below is the administrator's view of the dashboard. It includes information about:

- topology in which this VCDA installation participates – how many pairings with cloud and on-premises instances are configured
- brief health report for the most critical components participating in solution operations
- number of protected workloads grouped by VMs or vApps, direction and remote site type (cloud or on-premises)
- recent tasks and their result
- resources (CPU, memory and disk space) required to failover incoming replications aggregated and per tenant

Tenant's view of the dashboard is reduced to number of incoming and outgoing replications, consumed disk space and network bandwidth utilization with selectable interval and periods, information what are pairings to the specific tenant Organization and resources required to failover all incoming replications.

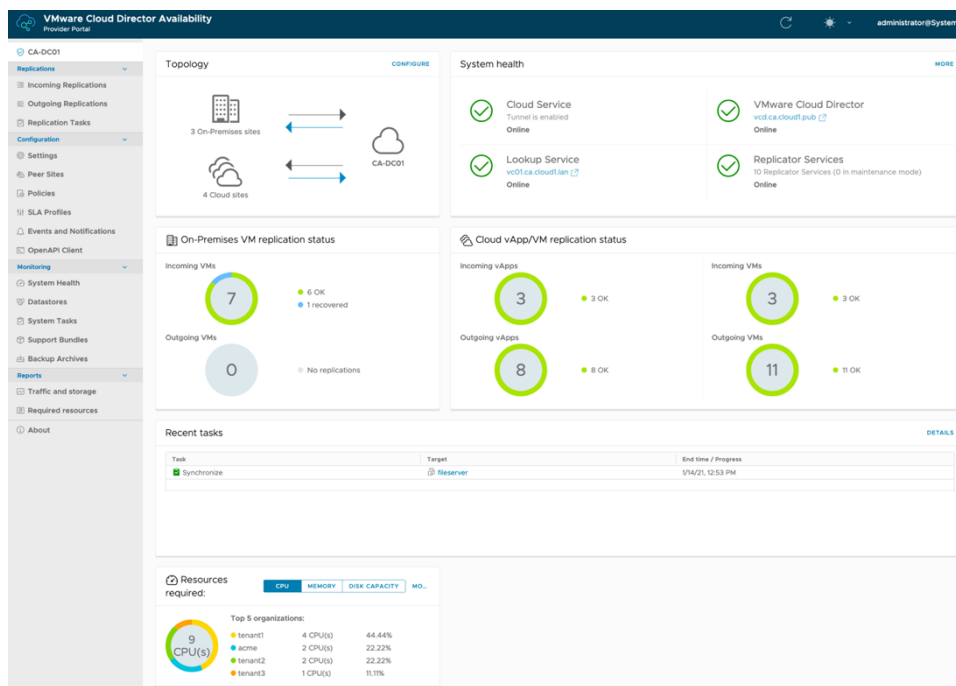


Figure 3 - VCDA Dashboard

System monitoring provides more detailed information about the health status of each local component which is part of the solution plus other components on which VCDA depends.

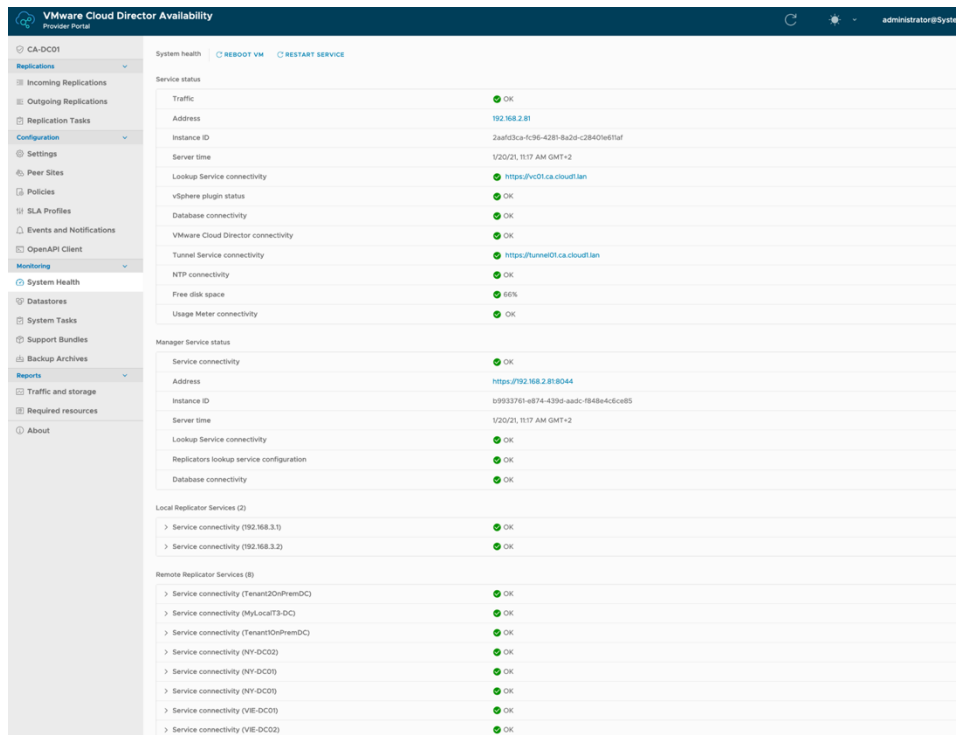


Figure 4 - VCDA System Monitoring

Other pages provide much more detailed information for usage. This is explained in detail in *Monitoring and reporting*.

VCDA API

VCDA provides a limited set of public APIs which is constantly extended with new API calls. This enables cloud providers to integrate VCDA solution with third party systems for different purposes. Polling usage information and monitoring are the most frequently used cases.

Information which are the public APIs in each release can be found in two places:

- VCDA portal Open API Client
- <https://code.vmware.com/apis/1016/vmware-cloud-director-availability>

VCDA extension in VCD Portal

VCDA can be integrated into VCD HTML Portal as an extension.

Note: There is no integration with Flash-based VCD interface.

When a user logs into VCD portal it is assigned an authentication cookie. When an organization is assigned to a VCDA Policy which has incoming or outgoing replications/migrations enabled, VCDA extension is enabled for any user that belongs to this organization.

If cloud provider did not configured VCDA-based DR service or if particular tenant is not enrolled to use this service VCD UI has an “Availability” placeholder which informs for the option to use VCDA to build a DR service.

VCDA extension is located in:

- hamburger menu for VCD versions before 10.1

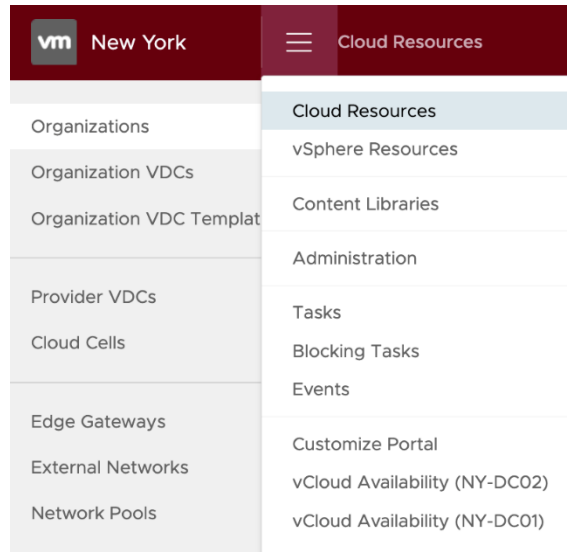


Figure 5 - VCDA extension in VCD 10.0.x and earlier

- under “More” in main menu in VCD 10.1 and later




Figure 6 - VCDA extension in VCD 10.1 and later

Only members of System organization and OrgAdmin users are allowed to load VCDA extension in VMware Cloud Director. When a user navigates to VCDA extension, the authentication cookie is reused and based on permissions assigned to logged in user VCDA portal is loaded into VCD Portal window. System organization members load the full VCDA UI while tenants OrgAdmins load UI with reduced functionalities and they can configure and operate replications/migrations and also get reporting limited to their Organization scope.

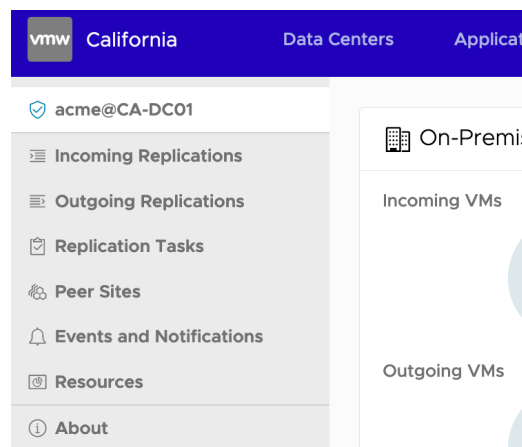


Figure 7 - VCDA as a VCD extension for non-administrator user

The VCDA portal is loaded into the VCD UI. It provides the exact same features. There is only one exception to this rule - tenants that have their own Identity Management service (e.g. dedicated SAML server) cannot log into the VCDA portal directly - they must use the VCD plugin.

Important details for successful integration between VCD and VCDA are explained in *Registering VCDA extension in VMware Cloud Director*.

VMware Cloud Availability Data Path

Two components realize replication data path – replicator and tunnel. Manager VM does not handle replication data except if combined appliance is used but it was earlier mentioned combined appliance is not supported for production installations and should be used only for tests.

Tunnel

Tunnel is a transparent proxy which is aware of TLS protocol. It receives incoming traffic and redirects it to backend components based on its type. Management and HTTPS/API traffic is sent to manager, replication data traffic is sent to replicator. It does not perform any operations on replication traffic.

Tunnel appliance does not consume significant compute resources but in environments with many replications may handle significant amount of network traffic. As it is published to Internet via DNAT it's recommended to connect it to a DMZ network where all security measures for such publish rules are taken.

Replicator

Replicator appliance runs three services – HBRsrv, LWDproxy and replicator.

HBRsrv is a low-level component which is communicating with replication IO filter in ESXi hosts. Protocol used for communication is called LWD – Light Weight Delta, and it's a VMware proprietary protocol. This service is not aware of vCenter and does not communicate with it. HBRsrv service is responsible to save VM deltas generated during RPO window on the destination datastore. HBRsrv is a common service for other VMware replication products and VCDA reuses it for moving replication data at the last mile.

VCDA configures the ESXi filter to send the replication traffic between ESXi host and source replicator in "plain text". When data is received by the source LWDproxy it is encrypted and optionally compressed. On the destination site there's another LWDproxy that decrypts and decompresses the LWD traffic and replication data is sent by destination replicator to an ESXi host.

LWDproxy is a VCDA service developed to take care to resolve this limitation and ensures replication traffic sent to the tunnel and outside of the local infrastructure is encrypted and compressed. LWDProxy service communicates with the tunnel for sending/receiving replication data. Encryption and compression result in increased CPU utilization in replicator appliance. This may reduce the number of replications a single replicator can handle.

Replicator service is a VCDA management service which sends configuration instructions to HBRsrv and LWDproxy services for every single replication that is handled by the replicator appliance. It understands VC objects like hosts, VMs, datastores, etc. In this way replicator discovers and prepares source VMs for replication and also creates objects (independent disks) for incoming replications.

The way these three services operate is important for decisions where to deploy and connect replicators.

In a cloud environment, it is supposed there will be hundreds or thousands of VM replications/migrations. Even if the number of replications is lower than the maximum supported replication handled by a single replicator, it's always recommended to have two or more replicators deployed for many reasons:

- With several replicators it will be possible to spread the load generated by replication over multiple hosts.
- Also, it will be possible to put replication traffic on more physical uplinks and distribute the load better.
- It will be also possible during maintenance window to put a single replicator in maintenance mode without affecting replication traffic.

Usually cloud providers separate management and resource vCenters/clusters. The number of hosts in resource vCenters/clusters is significantly higher than the number of hosts in management cluster. The next recommendation is to

deploy replicators on resource hosts and not in management cluster. The reason is it will be possible to deploy more replicators and create a DRS rule to keep replicator VMs on different hosts for better load distribution. Also, replication traffic path from replicator appliances to the replication network on the resource hosts will be enhanced.

Considerations for connectivity requires to take a look at two perspectives – communication between replicator and hosts and communication between replicator and tunnel.

Hosts may use management vmkernel interface to communicate with the replicator or it's possible to have a dedicated vmkernel interface for replication.

Using management vmkernel simplifies the configuration but significantly reduces the control options available to the administrator. Also, it provides a risk to route uncompressed replication traffic which is highly non desired.

It's recommended to use a dedicated vmkernel for replication traffic. The reason is in this way administrator will have a better control over infrastructure. Using NIOC, the administrator will be able to set shares for different types of vmkernel traffic. This enables to carry replication traffic over dedicated uplinks. Procedure to configure replication via dedicated replication vmkernel interface is:

1. Create a replication VLAN
2. Create a routing interface for this VLAN. It could be in physical infrastructure or a NSX edge.
3. Create a virtual port group for replication traffic and tag it with VLAN ID created in step (1).
4. Create a vmkernel interface, connect it to port group from step (3), configure it with an IP address from routing interface subnet. Do not configure a gateway for this interface
5. Enable “vSphere Replication” and “vSphere NFC Replication” tags on this interface.
6. Connect replicator appliances to port group created in step (3)
7. Configure replicators with IP addresses in the same IP subnet and configure the routing interface IP address as the gateway address.

This will ensure LWD traffic between hosts and replicators will always stay in the same Layer 2 domain and only encrypted and compressed replication traffic will be sent out by replicators via a Layer 3 routing device to the tunnel.

MultiNIC appliances

MultiNIC appliances are used to optimize replication traffic flow. Each VCDA appliance can be configured with more than one network interface but usually this is done on tunnels and replicators.

Configuring an additional network interface in VCDA appliance is required in two major cases:

- To bypass a routing device to prevent it from handling huge replication traffic – this is a scenario when communication between tunnel and local replicators needs to be optimized.
- When pairing session from remote sites are coming via not connected networks – for example cloud provider wants to use different interfaces in the tunnel appliance for pairing over public and private networks.

There is one important factor when VCDA tunnel is configured with more than one network interfaces. Tunnel can communicate with the other VCDA components (manager and replicators) in its local site only via one of its interfaces. When configuring the tunnel with more than one interface, this constraint has to be accounted always. Communication with remote sites can be handled by each network interface including the one used for communication with local manager and replicators.

There is no such constraint when a replicator or a manager has to be configured with additional network interface. Even though it's strongly recommended to avoid splitting communication between VCDA components in a local site over multiple interfaces.

Configuring additional interface in an appliance brings up additional question about routing and which interface to be configured with the default gateway. In most cases, default gateway has to be configured on the interface connected to Internet. If that's not the first interface named ens160, it will be required to move the default gateway to a different interface. Network interfaces can be configured using appliance UI or /opt/vmware/h4/bin/net.py tool (more information how to use this

python script is provided in *Network configuration*). When routing has to be reconfigured, the UI may not be the right choice and configuration should be done through appliance VM console.

Another operation that may be required is to explicitly configure VCDA service with which interface to use. The process behind VCDA service will listen on each interface in VM but application logic will choose the first one. This is usually the interface which is configured during OVF deployment. In case that the service has to use another interface, the service will require reconfiguration using CLI. The required configuration must be done under VMware Support guidance.



Please do not manually edit configuration files in any VCDA appliance. Use the management UI, API or CLI when configuration change is required.

Relations of other infrastructure components

Platform Services Controller

All VCDA appliances may be configured with a Lookup Service for two purposes:

- Service authentication and resource discovery – replicators and manager require registration with Lookup Service for the purpose to authenticate to resource vCenters and discover resources which is required for operations from VCDA workflow.
- Authentication to VCDA service by administrator – this is not mandatory. This is the reason for the dotted line in Figure 2 - Network flow in VMware Cloud Director Availability between tunnel and VC/PSC.

It's important to mention that Lookup Service in the question is the one used by vCenters providing compute resources to VMware Cloud Director. If for example management and resource vCenters have different Lookup Services and VCDA manager is hosted in management vCenter, VCDA manager will have to be configured with the Lookup Service of resource vCenters from/to which it will protect VMs/vApps.

While discussing Lookup Service another key factor to mention is that a Lookup Service and associated Single Sign-On (SSO) domain is the boundary of a single VCDA instance. Because a single VCDA manager can be registered with a single Lookup Service, this determines that this VCDA manager and its related VCDA components will be able to provide DR service only for workloads in this SSO domain. This means if cloud provider has PVDCs in more than one SSO domain it will be required to deploy a VCDA instance (tunnel, manager, replicators) per each SSO domain.

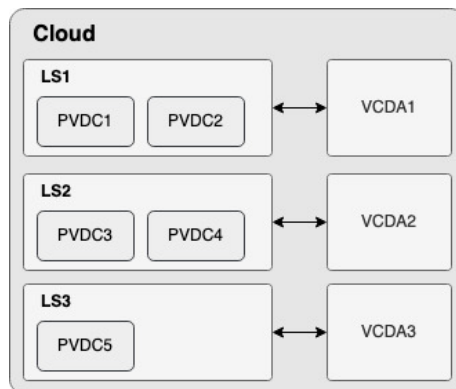


Figure 8 - VCDA to Lookup Service relation

It's possible to deploy VCDA instance per PVDC but a single VCDA instance can protect as many PVDCs exist in a single SSO domain.

vCenter Server

VCDA needs to be able to discover workload VMs that need to be protected or migrated and also destination resources (clusters, datastores, networks, etc.) for incoming replications and/or migrations in vCenters that are configured to provide resources in VMware Cloud Director.

Replicators are the components that communicate with vCenter Servers. Tunnel has no such requirement.

VMware Cloud Director

Cloud service in VCDA manager is communicating with VCD to discover vApps/VMs for outgoing replications/migrations or destination resources (OrgVDC, network, storage) for incoming replications/migrations. VCDA uses VCD API to request information and to initiate operations in VCD.

Integration between VCDA and VCD happens from vApp Replication Manager. VCDA expects to receive VCD API HTTPS Base URL followed by `/api` and user account which is a member of System organization with its password. After that VCDA inspects certificate presented by VCD and expects to find a FQDN of VCD API HTTPS Base URL as a Subject Alternative Name DNS entry. If certificate is not prepared in this way VCDA rejects to complete registration with VCD. It's not enough to have this FQDN present in certificate Common Name field. Once VCDA is able to establish a TLS connection to VCD & login with the provided credentials, VCDA uses VCD's REST API to discover its UI endpoint. Depending on its value, the VCDA HTTP server is dynamically reconfigured - the respective domain is put in the CORS whitelist.

Additional Information: *CORS is an HTTP mechanism that provides better security to end users. Depending on how the server is configured, browsers may reject attempts to load resources from a server that is on domain A, but someone tries to load it in domain B. The VCDA cloud service is configured to allow only one vCD endpoint - that is the vCD UI endpoint or if it's not configured, the API endpoint. So, for example, if the vCD UI endpoint is "ui.vcd.com", but you open vCD from "vcd.com" and try to open the VCDA plugin, your request is going to be rejected.*

Once certificate and credentials check are successful VCDA deploys in VCD cells files required for extension operation.

ESXi hosts

Replicators communicate with hosts to:

- Configure source VM for replication by enabling IO filter and configure it where to send replication data
- Allocate resources in host for destination replication

These operations are done by HBRsrv service. For authentication purposes it needs to talk to hosts on 80/tcp, for data traffic it uses 902/tcp to hosts to send replication data for incoming replication and 44046/tcp from hosts to receive replication data for outgoing replications.

Initial service configurations

Initial configuration of VMware Cloud Director Availability has changed in VCDA 4.1. Versions before 4.1 required to login on each appliance and conduct specific configuration. In version 4.1, configuration is much simplified and does not require to walk through each appliance. Both ways to configure VCDA are explained below.

Each VCDA appliance is deployed from OVA file. During OVF deployment administrator provides information where appliance will be deployed and connected in the hosting environment plus information for appliance guest OS configuration for:

1. Role of the appliance – manager, replicator, tunnel or combined appliances
2. Hostname
3. Search domain
4. IP configuration of network interface. Appliance is deployed with a single network interface. If additional interfaces are planned to be used, they are configured after deployment.
5. Initial password
6. Time server
7. DNS server

After an appliance is deployed and powered on, it configures OS settings necessary for appliance operations and then loads services based on the role selected during deployment. During its first load each service creates a self-signed certificate with validity of 1 year.

VCDA 4.1 and newer

If you have chosen to build your DR solution with VCDA 4.1 and newer, you will be able to use Initial configuration wizard from cloud manager after you change the root password at first login. In order to use it, it is required to deploy all VCDA appliances in your cloud and to have all network prerequisites in place. After that, all configurations are performed from vApp Replication Manager UI using Initial Configuration Wizard.

This wizard will request all required information to configure VCDA-based DR solution:

- License key
- VCDA Site Name and Public Service Endpoint
- VMware Cloud Director Public HTTPS URL together with System administrator credentials
- Lookup Service URL and credentials for it
- One or more replicators URL with credentials – if you have not logged into replicators before that, the wizard will recognize root password has to be changed and will ask you for a new root password
- Tunnel URL – if you have not logged in earlier on the tunnel appliance, the wizard will ask you for a new password and will replace the one provided during OVF deployment.

Note: Tunnel appliance is not configured with Lookup Service by Initial Configuration Wizard.

Once all these details are provided, your VCDA instance will be fully configured and ready for operations. To enable incoming/outgoing replications/migrations, change the Default Policy or by create a new policy and assign tenants to it.

VCDA 4.0.x and earlier

If you have still decided to use a version prior to version 4.1, it's required to prepare each appliance before combining all appliances into a single solution.

Replicator

To start initial configuration of replicator appliance, the administrator has to login to https://replicator_ip_or_fqdn:443 using root account and password provided during OVF deployment. Right after successful login, the appliance will request to change the root password. This is a security measure to remove the chance to reveal root password from OVF properties where password is saved in plain text.

Next step is to register replicator with the Lookup Service. It's important to use Lookup Service used by vCenters which will be source or destination for replications/migrations. If a replicator is deployed in different SSO domain than the one where VMs that will be replicated are hosted, the replicator still has to be registered with the workload Lookup Service.

Replicator Manager

To start initial configuration administrator has to login with root account using the password provided during appliance deployment on https://manager_ip_or_fqdn:8441. As always, the first step will be to change this password with new strong password.

Second step is to register manager service with the workload Lookup Service. The same rule for Lookup Service applies to this service also.

The third step is to register replicators in the manager. This step has to be performed after cloud service configuration is completed. The reason is that during replicator registration replicator has to be assigned to local site and site is created during cloud service configuration.

Tunnel

Tunnel service management interface is on https://tunnel_ip_or_fqdn:8442. First login to this interface has to be done with root and password provided in OVF deployment. Again, the administrator will be asked to change this password before doing anything else with the appliance.

Registration with Lookup Service is optional for the tunnel. It's required only if login to tunnel management interface has to be enabled for users from an SSO domain.

vApp Replication Manager

Configuration of cloud service completes the DR solution based on VMware Cloud Director Availability. It ties together configurations of other services and integrates into VMware Cloud Director HTML Portal.

Configuration page of cloud service is available on https://manager_ip_or_fqdn:443. If the administrator has not changed the appliance root password earlier by login into manager service interface, on the first login to cloud service UI change password window will pop up. If appliance root password has been changed earlier, the administrator can continue with Initial wizard which guides to steps required to complete service configuration. Initial wizard here does not configure any other service but vApp Replication Manager.

During this wizard the following items are configured:

- Site Name – It is very important to plan properly what the site name will be. Once it is set, it cannot be changed. If site name has to be changed, the only option is to redeploy manager appliance, and this could be a real problem if there are already replications configured.
- Lookup Service – this is again the lookup service for workload domain. Even if manager is hosted in a management vCenter which does not share the same SSO domain with workload vCenters, cloud service has to be registered with the Lookup Service of the workload SSO domain.
- Service endpoint information – this is the public URL that will be used by customers of the DR service. This setting is critical for integration with VCD. URL provided at this step will be used by VCD to call VCDA extension in its HTML interface. Usually, this URL is resolved, by external users, to a public IP address through which VCDA is accessible from Internet.
- VMware Cloud Director URL – this is the VCD URL from which VCDA will accept to be loaded as an extension into VCD HTML portal. VCD way to control this is through its setting “Public API Base HTTPS URL”. That's why it's recommended to provide the same URL at this step. As already mentioned in VMware Cloud Director section, VCDA has a requirement about how VCD certificate is prepared. VCDA expects that VCD certificate has a SAN DNS entry matching VCD Public Base HTTPS URL.

With this the initial wizard is completed and there are two additional steps required to have a VCDA site capable to provide DR service – register local replicators to manager service and enable tunneling which is a mandatory step.

To register one or more local replicators administrator has to load manager service URL on port 8441 and from **Replicator Services – New**, start the wizard for registration of new local replicator.

Tunnel enablement is explained later in this chapter.

Using wizards helps a lot to prepare configuration but hides some details. In the following sections, the most important configuration settings are explained in more details.

Public API Endpoint

In VCDA versions prior to version 3.5, this setting had much more important role. It was not possible to pair to anything that does not match what has been configured for public API endpoint. Since version 3.5 and later, it is possible to pair to any URL which reaches the tunnel on port 8048/tcp.

The Public API Endpoint is important because when VCDA is registered with VMware Cloud Director, this URL becomes the URL that VCD will call when it will load VCDA as an extension.

If VCDA Public API endpoint URL has to be changed, registration with VCD has to be re-established to update Cloud Director about the change in VCDA.

Registering VCDA extension in VMware Cloud Director

When VCD is registered with VCDA, VCD registers an UI plugin that points to VCDA Public API endpoint. At the same time, VCDA adds VCD Public HTTPS URL in its allowed CORS list and configures itself to accept requests to load inside VCD browser frame only when it's called from URL that matches VCD Public HTTPS URL. If VCD Public HTTPS URL has been changed and VCDA is not aware of this change VCDA will reject to load into VCD browser window.

This is explained in detail in *VMware Cloud Director*.

In case of changing VCD Public HTTPS URL or renewal of certificate bind to VCD Public HTTPS URL, it is required to re-register VMware Cloud Director with VCDA. This will update URL together with certificate information from which VCDA will accept to be loaded as an extension.

In all cases, when VCD is registered with VCDA, VCDA will check the certificate presented by VCD. VCDA will require that the FQDN configured as VCD Public HTTPS URL to be present as a Subject Alternative Name DNS entry in this certificate. If this requirement is not met, VCDA will reject to complete such registration.

Enable Tunneling

As already mentioned, tunnel is a mandatory VCDA component and all incoming and outgoing replication traffic should pass through the tunnel appliance. To achieve this, tunneling has to be manually enabled. When tunneling is enabled, all local components are prepared to communicate with the tunnel.

All VCDA components talk to each other through a secure connection that's established in a TLSv1.2 session. Whenever component from site A wants to talk to component in site B, the component A uses the TLS SNI (Server Name Indication) extension. The tunnel knows how to read the ClientHello message (which contains the SNI) and determines where to send the traffic (the site B component endpoint) based on its internal routing table. Note that this does not require multiple TLS sessions - the tunnel does NOT store private keys and does no TLS termination. It is a simple TCP proxy that forwards packets based on the TLS SNI extension. This is one of the many reasons why VCDA requires end-to-end encryption and does not support having products that perform TLS MITM.

Enable tunneling is performed from cloud manager UI. Navigate to Settings – Service Endpoints – Tunnel Service Address – Edit and in the window that will pop up the following has to be provided:

- Tunnel appliance IP or FQDN with port 8047
- Root account – it is auto-populated
- Appliance root password

Manual tunnel enablement may be required in case of:

- Renewal of vApp Replication Manager certificate
- Tunnel service certificate renewal
- Re-deployment of tunnel appliance

It's important to note the purpose of the two tunnel ports. 8047/tcp is tunnel API port used for management traffic between local components and the tunnel. That's the port use to enable tunneling. Port 8048/tcp is the data endpoint used by tunnel to receive traffic from remote tunnels and on-premises replicators.

SSO Registration

As explained earlier in this document, vApp Replication Manager has to be configured with Lookup Service for resource vCenter and not for the management vCenter. This registration is required to enable users to use SSO credentials when log into VCDA portal.

If Lookup Service certificate is replaced, it will be required to re-establish the channel between Lookup Service and vApp Replication Manager service, Replicator Manager Service and all replicators.

Pairing

General information

Pairing is the process of establishing a trust between two VCDA clouds or between an on-premises replicator and VCDA in the cloud.

To make it possible to configure a replication, the user needs to establish login both in the local site and in the remote site. Login to remote site happens over the trust established with pairing.

This trust channel together with authenticated sessions enable discovery of:

- workloads to be protected or migrated
- destination resources for protections or migrations

Pairing between clouds is a two-step process and every step is performed from different cloud sites. Pairing between clouds is an administrator operation. To complete pairing between clouds, the administrator initiates pairing and provides remote site name and URL.

The components use X509 certificates for authorization & authentication. In order to complete the pairing, the administrator is prompted to verify the certificate of the remote component. After administrator accepts the certificate, a secure connection is established. Administrator is prompted to accept the certificate assigned at the remote endpoints.

Pairing between on-premises and cloud is initiated from the on-premises site. The on-premises administrator provides on-premises site name, VCDA cloud URL and credentials from an VCD Organization to which on-premises will replicate workloads and this completes the process. Pairing wizard provides an option to enable or disable access to on-premises vSphere inventory from cloud. There are two cases when this has to be enabled:

1. When the cloud provider and the tenant agree that the provider will provide managed services on behalf of the tenant. In this case, the tenant does not need to provide access to its vCenter directly to the cloud provider, but the provider will be able to use VCDA portal or VCDA extension in VCD to browse only VMs from on-premises datacenter that has to be replicated or migrated to the cloud.
2. When the on-premises vCenter is running an old version that does not have HTML5 client. In such cases, using VCDA portal or VCDA extension in VCD is the only method to use VCDA to replicate or migrate the required VMs.

From cloud standpoint, when pairing happens, Replicator Manager service registers in its replicators inventory all replicators from remote site. This enables it to create pairs of local and remote replicators for incoming replications for which the local Replicator Manager is responsible.

A single tunnel is capable to accept pairing sessions either over private or public network. Details about each of these are in the next two sections.

Pairing over Internet

Pairing over Internet usually happens to URL configured as Public API endpoint but this is not mandatory. It's enough to use URL that resolves to network device which does DNAT to the tunnel appliance and it's required to configure DNAT in a way that traffic is delivered to tunnel interface on port 8048/tcp. Public port is not important, only translation port is important. If cloud provider plans to publish VCDA via IP address for which 443/tcp is already used to publish other service, it's fine to use any other tcp port but this port has to be translated to https://tunnel_IP_address:8048.

Pairing over private networks

Pairing over private networks does not go through any DNAT device in front of the tunnel. It needs only to reach any tunnel interface on port 8048. It's either possible to use IP address of any tunnel interface or FQDN which resolves to tunnel IP address.

Admin access from Internet

This is a security setting which controls if administrative sessions will be authenticated or not when they originate from public network. If enabled, this restriction applies to the following:

- Login sessions of local root account
- Login sessions of VCD system administrator
- Login sessions of SSO Domain Administrator

This does not mean web page `/ui/admin` won't be loaded. It means if it's disabled any login attempt will be unsuccessful and will result in `401 Not Authenticated`.

Throttling

There are two options to enforce bandwidth throttling:

- how much replication traffic will enter the local site

- how much traffic an on-premises replicator will send to the cloud

The first option is from VCDA Portal Settings page. This global setting controls what the allowed amount of incoming traffic that tunnel will handle from all remote sites and this limit is applied on the interface used by tunnel to communicate with the other VCDA appliances in the local site. In this way, VCDA prevents scenario when storage and network infrastructure in local site is heavily utilized by replication traffic.

Figure 9 - Bandwidth Throttling Configuration

The second option is to control the amount of traffic that an on-premises replicator will send to the cloud and this is possible using policies. This throttling setting in policies is applicable only for on-premises replicators and not for cloud replicators. When bandwidth throttling is set, the on-premises replicator respects this limit and tries to optimize its operations in a way that RPOs will be met without violating the bandwidth cap. If that's not possible due to a low bandwidth limit, large deltas or short RPO windows, then RPO violations will be observed and it will be required to optimize the settings by changing allowed bandwidth, RPO windows or replication settings.

Policies

Policies control what operations are allowed for tenants assigned to the policy and enforces some limitations. The settings that could be managed by a policy allow cloud provider to keep control over utilization and performance demand generated by replications and migrations. Policies are under control of cloud provider and tenants use them as configured by the cloud admin.

A policy distinguishes between replications and migrations and controls which direction is enabled or disabled. Also, it may set limits to:

- Maximum number of configured replications
- the minimal RPO window allowed for replications that tenant will configure
- maximum number of pinned and rotated MPITs
- Bandwidth limit for the on-premises replicator to limit the outgoing replication traffic

Using policies cloud provider may allow or forbid the tenant the option to forward VCDA events to VCD.

SLA profiles

Policies also control if tenant is allowed to use SLA profiles for its replications. SLA Profile consists of specific settings applicable to each replication. The idea behind SLA profile is to pre-configure these settings and store them in SLA profile for future use. The benefits are:

- A replication is configured with less clicks
- Each workload could be assigned with SLA profile and could have consistent protection settings
- Regular users could configure replications with settings usually managed only by admins

A single SLA Profile consists the settings below:

- Profile Name
- RPO – from minimum allowed by policy to 24h
- Retention policy – number of MPITs
- Enabled or disable Guest OS quiescing
- Enable or disable compression of replication traffic
- Delayed start of initial synchronization

When tenant is enabled to use SLA Profiles during replication configuration, the tenant may choose to use or not an SLA Profile. If SLA Profile is selected, replication settings that are required to be configured are significantly reduced as show on the picture below:

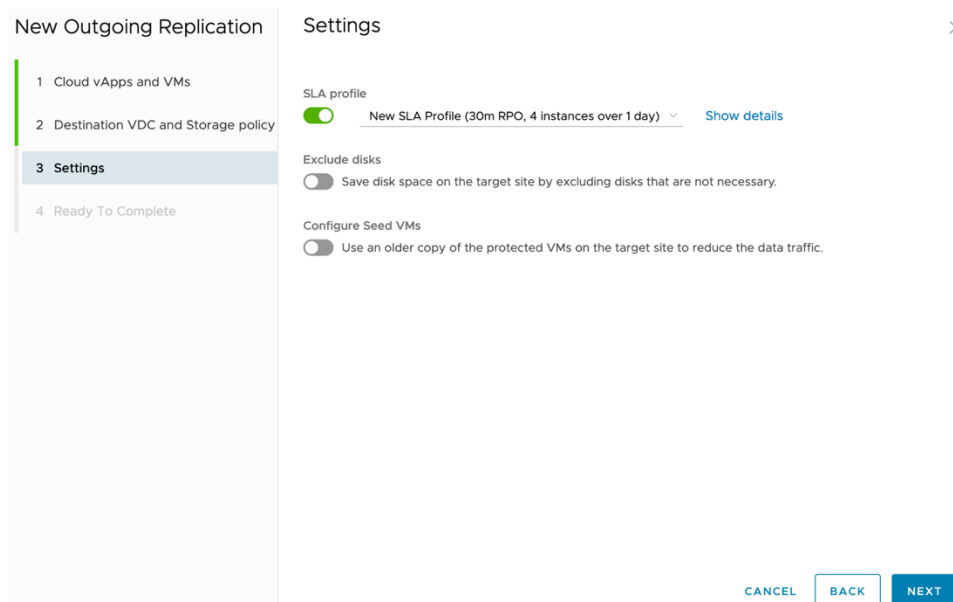


Figure 10 - Configure replication using SLA Profile

VMware Cloud Director Availability at on-premises site

Overview

One of the use cases of VMware Cloud Director Availability is protection and migration of workloads hosted in vSphere managed environments to the cloud. To make this possible VMware provides a dedicated VCDA appliance for on-premises deployment. This appliance is a replicator which is configured to protect and migrate to/from the cloud VMs managed by vCenter Server.

For modern vSphere environments VCDA on-premises appliance supports native integration with vSphere Client through a plug-in which is deployed in vCenter and it's a part of vSphere HTML5 Client. Such integration is not possible with legacy management tools like C# and Flash clients. Customers who have old infrastructures without HTML5 client can use VCDA Portal to configure and manage replications or migrations to/from on-premises site.

VCDA on-premises appliance supports replications and migrations in both directions – incoming from the cloud and outgoing to the cloud. Actual capabilities depend on policy settings assigned by cloud provider to the organization in the cloud.

VCDA on-premises appliance is shipped also in OVA format and is available for download from <https://my.vmware.com>. It should not be downloaded and distributed by the cloud providers for legal reasons. Only customers can download OVA after they accept EULA before using the product. Cloud providers should only provide instructions for versions they support.

Requirements

VCDA on-premises replicator requires to communicate with vCenter and hosts in the on-premises site. Exact endpoints, protocols and port can be found in documentation and also on <https://ports.vmware.com>. It's important to mention that if SSO domain is hosted on vCenter which has been upgraded from 5.5 Lookup Service listens on port 7444/tcp and when on-premises replicator is registered with Lookup Service it's required to provide explicitly the correct port in Lookup Service URL.

VCDA on-premises replicator does not accept any incoming sessions from Internet. This means there is no need to configure any service publishing from on-premises replicator. As it initiates pairing to VCDA in the cloud it's required for on-premises replicator to have outbound connectivity to the cloud via Internet or via private network.

Configuration

After deployment, the first task is to change the root password provided as part of OVF deployment.

Next configuration steps can be grouped in two operations: pairing configuration and local placement configuration.

Pairing

Pairing is the process of establishment connectivity and trust between on-premises replicator and VMware Cloud Director Availability in the cloud. Before pairing is initiated, the cloud provider has to provide its tenant with the following details:

- VCDA public endpoint URL
- VCD Organization admin username
- Password
- Certificate thumbprint if VCDA cloud service does not use a certificate signed by well-known public certificate authority.

A dedicated wizard helps an administrator to complete pairing process. The following details has to be provided on several steps:

- On-premises site name – this will be visible in cloud provider on-premises site lists and will distinguish one from another on-premises tenant
- Lookup Service URL and credentials – on-premises replicator registers itself into Lookup Service. Using this registration, vCenter discovers the on-premises replicator and downloads from it the plug-in which integrates into vSphere HTML Client.
- Cloud Service details – at this step, administrator provides information for VCDA installation in the cloud which includes VCDA cloud service endpoint, VCD organization admin username and password. For successful pairing, it's required that VCD Organization is assigned to a VCDA policy with enabled replication.

There is an important setting to be configured at this step - to allow or disable browsing on-premises vSphere inventory from the cloud. Two obvious cases when this could be enabled is when cloud provider will provide managed services and when vSphere version in on-premises does not support plug-in integration with vSphere Client.

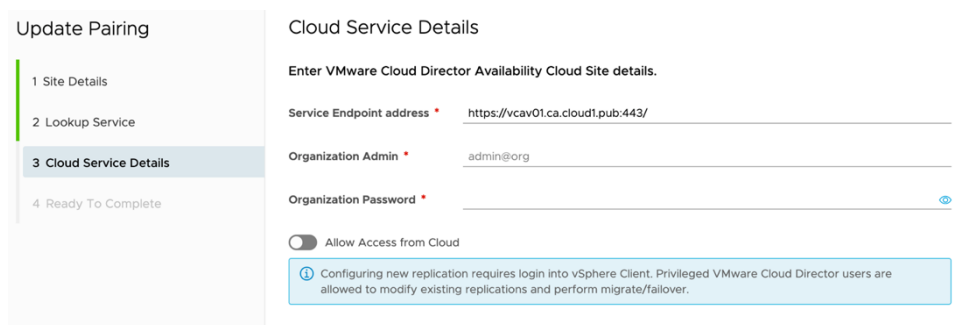


Figure 11 - Allow Access from Cloud

Once administrator presses **Next**, the Administrator is presented with certificate thumbprint of the VCDA service in the cloud. After validation of the certificate and acceptance, the wizard reached its final review window. On this window, the administrator may want to enable or disable the switch which controls if placement wizard will start right after pairing wizard closes.

Placement

As already mentioned, VMs could be failed over or migrated from the cloud to the on-premises infrastructure. Placement settings define where such VMs will be hosted, connected and stored in on-premises vCenter inventory. Placement wizard collects information regarding folders in which VMs will be put, clusters and resource pools where VMs will be hosted, datastores VMs files will be stored and networks where VMs network interfaces will be connected.

Appliance management

Majority of appliance settings are configured during deployment by OVF properties. Some of the settings are used as they are configured during deployment; others may be changed later. Passwords provided as OVF properties have to be changed after first login.

Network configuration

Network settings of an appliance can be managed in two ways – via UI and via a Python script `/opt/vmware/h4/bin/net.py`. Manual editing of network scripts in `/etc/systemd/network` directory is not recommended and supported unless this is a recommendation from engineering team.

If during deployment wrong OVF properties were provided appliance may not start with fully configured network interface. In such cases, the Python script can be used to fix the error and after that UI can be used for any network configuration change that may be required. The operations supported by this Python script are:

- `configure-nic` - updates the configuration of a specific nic
- `nics-status` - displays the current configuration of all nics
- `nic-status` - displays the current configuration of a specific nic
- `unconfigure-nic` - unconfigures the specific nic and makes it unroutable
- `dns-status` - displays the current DNS configuration
- `configure-dns` - changes the global DNS configuration
- `list-routes` - retrieves the currently configured static routes
- `add-route` - creates a new static route
- `remove-route` - deletes a static route

All these operations are available through UI too. Both UI and python script can be used to configure any network interface in the appliance.

NTP

VCDA appliances are required to be synchronized to the same NTP server(s) used by all other components (vCenters, VCD cells, hosts). NTP source can be provided with OVF property but it can be also configured later through UI or python script `/opt/vmware/h4/bin/configure-ntp.py`. If an appliance can communicate with NTP source and can get time updates from it, Dashboard will report this with a green mark. If not, Dashboard will have an alert, and this has to be fixed.

Failure to synchronize time may result in failed authentication and communication between VCDA components. Hence, it's critical to ensure NTP time sync is working smoothly.

Certificate Management

Certificates are critical part for VCDA operations. Each appliance authenticates to other appliances using its certificate. VCDA appliances can use equally successfully self-signed certificates and certificates signed by Certificate Authority. On its first load, each service is configured with a self-signed certificate which is valid for 365 days.

Certificate replacement can be done through UI.

Appliance settings		
Root password	*****	Change
> Network	vcav01.ca.cloud1.lan 🔗	Edit
▼ Certificate	4/21/20, 10:03 PM - 4/21/21, 10:03 PM	Import Regenerate
▼ Issued To		
Common Name (CN)	cloud.vm	
Organization (O)	Unknown	
Organization Unit (OU)	Unknown	
▼ Issued By		
Common Name (CN)	cloud.vm	
Organization (O)	Unknown	
Organization Unit (OU)	Unknown	
▼ Fingerprints		
SHA-256 Thumbprint	SHA-256:2A:EA:3D:4E:57:C5:E4:23:47:B7:34:0C:8C:C0:B6:81:B2:D3:65:F9:FE:64:36:56:C8:96:65:AC:37:6B:EF:8E	

Figure 12 - Certificate renewal in VCDA

Using Import option, CA-signed certificate can be assigned to the service. CA-signed certificate has to be in PKCS#12 format. Regenerate option will generate a new self-signed certificate with 365 days validity. After new certificate is imported/regenerated service is automatically restarted to load with the new certificate.

Key notes regarding certificates:

1. VCDA expects that each component will have its own certificate. Avoid using wildcard certificates on all services as this may result in a failed VCDA setup.
2. Only certificate of cloud service running in VCDA manager is visible to external world. This certificate is presented during pairing and when users load VCDA portal directly or as a VCD extension. If a signed certificate has to be used for your DRaaS by VCDA, this is the only service that needs such certificate. The rest of the services – tunnel, manager, replicators can continue to operate with self-signed certificates.
3. Do not install certificate on any security device which may exist in front of the tunnel. SSL operations like certificate replacement or SSL offload will be recognized by VCDA as a traffic modification during transit and will result in rejection to authenticate and further processing. Most frequently this is observed as failure to pair sites. If security device has been deployed in front of the tunnel, configure it to operate in transparent mode for VCDA, so that traffic reaches the tunnel in the original form as prepared by remote tunnel or on-premises replicator.

In this [blog post](#) you can find more details about the certificate management of VMware Cloud Director Availability.

Monitoring and reporting

Built-in reporting

Information provided by built-in reporting tools is used by the cloud provider as an entry point for its cost model to generate monthly bill for its tenants.

Since version 4.0, VCDA Portal can provide reporting information in its UI to the cloud provider for traffic data and consumed storage per tenant for certain period of time. Information is provided in both graphical format and also in raw format if the provider needs to export it.

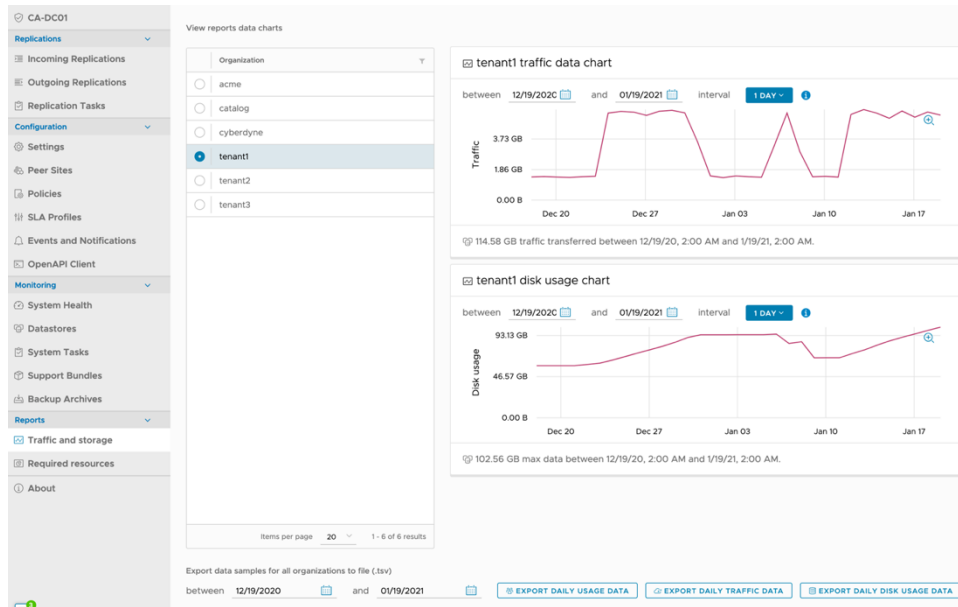


Figure 13 - VEDA portal reporting

Also, both cloud provider and tenants are provided with information about the amount of compute resources required to fail over or migrate workloads which are currently replicated to the cloud site by VEDA. This includes information about the number of virtual CPUs, memory and disk space used by all incoming replications. Below is the tenant view of this report:

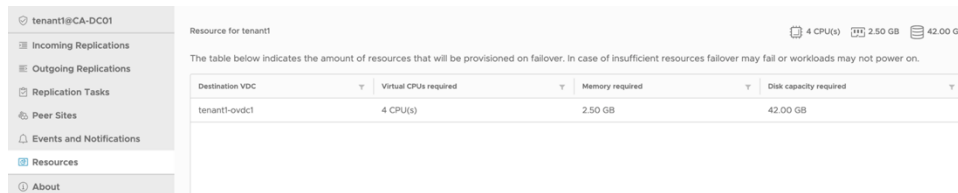


Figure 14 – Required resources per tenant

Cloud providers that still use vCloud Availability 3.x can pull this information only using CLI. Steps how to do this are described here:

<https://docs.vmware.com/en/VMware-vCloud-Availability/3.5/administering-vcloud-availability/GUID-DC847544-6214-4242-A3E2-D5B490DDE8A7.html>

UM integration

As part of VMware Cloud Provider Partners (VCPP) program, Usage Meter provides licensing usage metering and automated usage reporting for resources owned by cloud providers and managed with VMware products. As VEDA is accessible to cloud providers through VCPP program, it is tightly integrated with Usage Meter and eliminates the need for human activities to generate and submit monthly reports about what resources have been active and managed with products from VCPP program.

VEDA is compatible with all latest Usage Meter versions. If you have 3.6.1 please ensure it has hot patch 5 or later. Here is the link to VEDA/Usage Meter interoperability matrix:

https://www.vmware.com/resources/compatibility/sim/interop_matrix.php#interop&570=&662=

Information about how Usage Meter can be deployed in a cloud provider environment is available on Usage Meter documentation page. Once deployed and configured, Usage Meter needs to be integrated with all VEDA Managers from which Usage Meter will collect and report usage statistics. The link below is from Usage Meter 4.3 documentation and provides the steps to register VEDA Manager to Usage Meter:

<https://docs.vmware.com/en/vCloud-Usage-Meter/4.3/Using-and-Managing-vCloud-Usage-Meter/GUID-7C778B16-299C-477B-8D99-C3B753DA8042.html>

VRLI pack

vRealize Log Insight VCDA Content Pack is available from VRLI Marketplace. VCDA logs and events serve as a base for the build-in dashboards, queries and alerts which enable cloud provider to be informed how VCDA operates and if attention is required.

Cloud admin has to configure syslog server URL that has VCDA Content Pack installed. This can be configured from VCDA portal – Events and Notifications option. For deeper analysis, cloud admin may deploy VRLI agent to VCDA appliances.

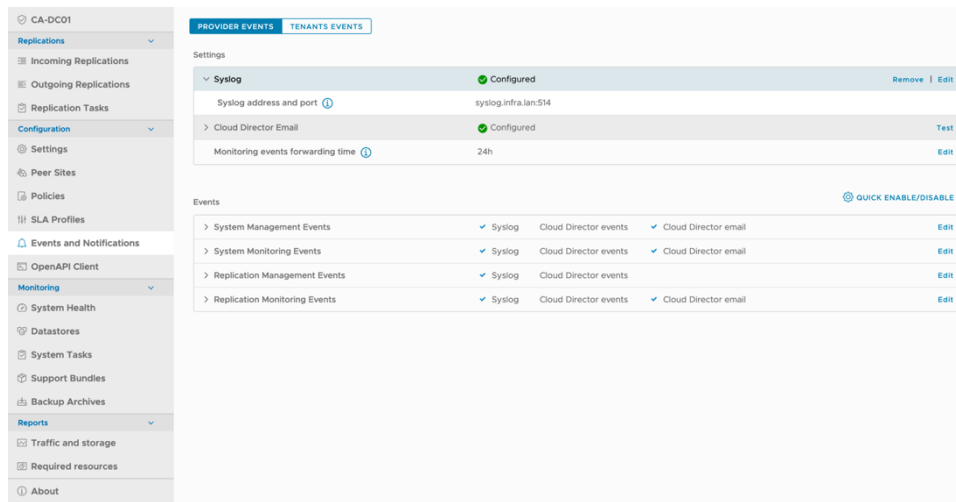


Figure 15 – Provider notifications configuration

VROPs pack

vRealize Operations is used in almost every cloud provider. VCDA vROPs Management Pack (VCDA MP for vROPs) provides detailed insight over VCDA protected workloads. It is available from <https://marketplace.cloud.vmware.com/services/details/c372290d-9997-4f85-b4d4-bfc1e5c71cc7>.

VCDA MP for vROPs is distributed as a single .pak file and provides both provider and tenant incoming replications view with detailed information about health states, configuration settings and consumed resources.

Screenshot below demonstrates how tenant can drill down into workloads protected with VCDA and get information for the smallest details:

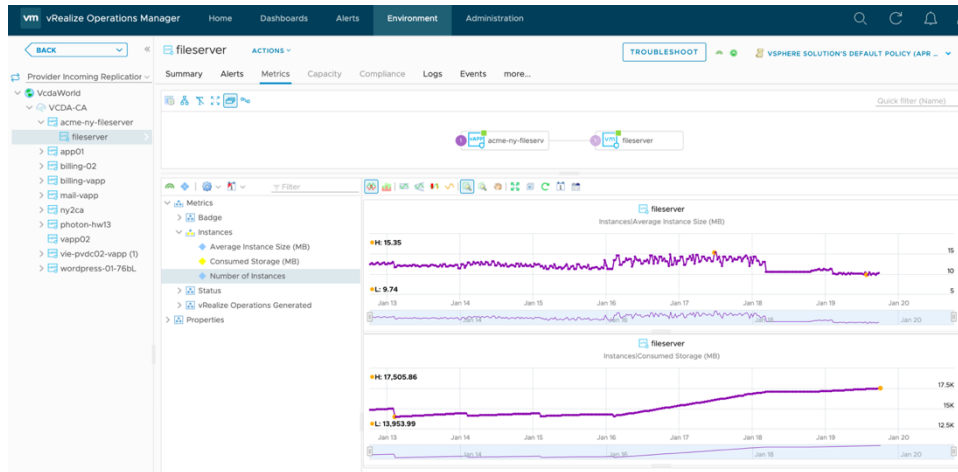


Figure 16 – vROPs reporting through the management pack

List of Figures

Figure 1 - VMware Cloud Director Availability Components 6

Figure 2 - Network flow in VMware Cloud Director Availability 8

Figure 3 - VCDA Dashboard10

Figure 4 - VCDA System Monitoring11

Figure 5 - VCDA extension in VCD 10.0.x and earlier12

Figure 6 - VCDA extension in VCD 10.1 and later12

Figure 7 - VCDA as a VCD extension for non-administrator user12

Figure 8 - VCDA to Lookup Service relation15

Figure 9 - Bandwidth Throttling Configuration21

Figure 10 - Configure replication using SLA Profile.....22

Figure 11 - Allow Access from Cloud23

Figure 12 - Certificate renewal in VCDA25

Figure 13 - VCDA portal reporting.....26

Figure 14 – Required resources per tenant.....26

Figure 15 – Provider notifications configuration27

Figure 16 – vROPs reporting through the management pack.....28

Reference Documents

Item	URL
VMware Cloud Availability Documentation	https://docs.vmware.com/en/VMware-Cloud-Director-Availability/index.html
VCDA Public API Documentation	https://code.vmware.com/apis/1091/vmware-cloud-director-availability
VCDA Network Ports	https://ports.vmware.com/home/VMware-Cloud-Director-Availability
Interoperability matrix	https://www.vmware.com/resources/compatibility/sim/interop_matrix.php#interop&570=4715,4716,4096,3622&1=4275,3495,3456,3221,2861,3363&2=4276,3496,3457,3222,2862,3364&93=4795,3608&175=4729,4232&224=3535&661=4254,4671,4249,4058,3535,3101&662=4748,3907,2666&783=4795,3608

Acknowledgement

Author

Atanas Stankov – Senior Solution Architect, CSBU

Reviewers

Shady Ali ElMalatawey - Staff Solutions Architect, VVD Engineering

Viktor Vasilev – Senior MTS, CSBU



vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.
Copyright © 2020 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-wp-temp-word 2/19