VMware Cloud Services Cloud Proxy

Understanding the VMware Cloud Services Cloud Proxy for VMware vRealize Automation Cloud and VMware vRealize Log Insight Cloud

This document applies to the three VMware Cloud services that comprise VMware vRealize® Automation Cloud™: VMware Cloud Assembly™, VMware Service Broker™, and VMware Code Stream™, and VMware vRealize® Log Insight Cloud™.

TECHNICAL PAPER

SEPTEMBER 2021

VERSION 1.7

Table of Contents

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

Revision History

DATE	VERSION	DESCRIPTION
July 16, 2019	1.0	Initial version.
July 22, 2019	1.1	Minor edit.
September 6, 2019	1.2	Product name changes.
September 26, 2019	1.3	Update to information about cloud proxy initiation.
October 8, 2019	1.4	Minor edit.
December 6, 2019	1.5	Title change.
September 9, 2020	1.6	Language cleanup.
September 15, 2021	1.7	Add whitelist section to FAQs.

Summary

The VMware Cloud service's cloud proxy connects VMware Cloud Services Platform (CSP) services, such as VMware Cloud Assembly and VMware Code Stream, to on-premises data centers. It is a virtual appliance delivered as an OVA that you deploy on an on-premises vCenter Server. The cloud proxy calls containers of specific agents for the various VMware Cloud services and supports data communication between the cloud provider and the managed vCenter environment. The cloud proxy enables lifecycle management (LCM), as well as data delivery and communications functionality. Within the cloud proxy, service agents gather data and use a data pipeline service that provides high throughput and low latency data delivery. It also controls channel communication between the cloud proxy that has been deployed to a vCenter Server on-premises and the VMware Cloud services. The cloud proxy was referred to as the remote data collector (RDC) in earlier iterations of VMware Cloud services.

Note: This document applies to the three VMware Cloud services that comprise VMware vRealize® Automation Cloud™: VMware Cloud Assembly™, VMware Service Broker™, and VMware Code Stream™, and VMware vRealize® Log Insight Cloud™.

What is the VMware Cloud services cloud proxy

A cloud proxy is a:

- Virtual appliance (VA) that is supplied as a downloadable OVA from the VMware Cloud service. The OVA must be deployed on a vCenter Server to create the VA. Deployment to an ESX server is not supported. The VA is comprised of several Docker containers. During VA deployment, the relevant agents are downloaded to the appliance.
- Cloud proxy service client that handles automatic download, configuration, and LCM

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

capabilities, such as upgrade, of service-specific containers. The service is also responsible for initial registration with the specific VMware Cloud service.

- Service agents, such as for VMware Cloud Assembly or vRealize Log Insight Cloud, that provide functionality elements to the VMware Cloud service. Examples include pull functions for metrics, push functions for logs, and provisioning functions for automation.
- Connection to the data pipeline service that connects to the VMware Cloud service. A data pipeline service client is embedded in each Docker container in the VA for use as the data pipeline and command channel. The data pipeline service handles connectivity with VMware Cloud services, as well as authentication and buffer/retry.

What does the cloud proxy do

A cloud proxy:

- Connects VMware Cloud services to on-premises networks.
- Ensures highly secure, bi-directional communication using OTK (One Time Key) and public/private key cryptography.
- Is highly scalable and designed to be the single solution for use by the supported VMware Cloud services. The same cloud proxy is available for your vRealize Automation Cloud and vRealize Log Insight Cloud services.
- Provides self-healing functionality that is facilitated by Docker restart and monitoring capabilities.
- Uses a single data channel to communicate, via a high-performance data pipeline, between VMware Cloud services and the cloud proxy.
- Is available for download from the VMware Cloud Assembly, VMware Service Broker, VMware Code Stream, and vRealize Log Insight Cloud user interface. Installation and deployment instructions are supplied in an on-screen wizard and in-product documentation at docs.vmware.com.

How does the cloud proxy support SaaS

The following integration diagram conveys the basic mechanics of how VMware Cloud services such as VMware Cloud Assembly and vRealize Log Insight Cloud are connected to the cloud proxy on a target vCenter Server.

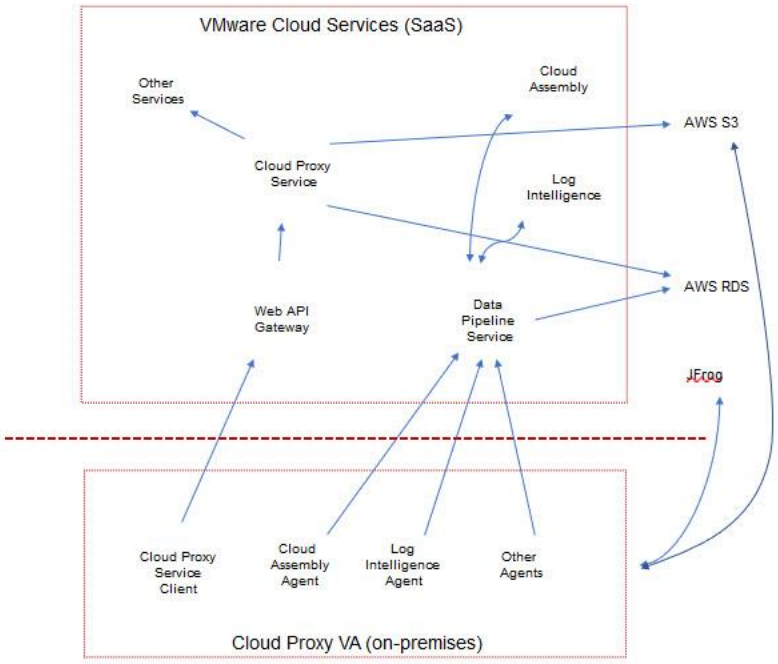

Common cloud proxy considerations are:

- **High Availability (HA):** You can use vSphere HA for the cloud proxy virtual appliance. HA is not available for agents running in the cloud proxy virtual appliance. However, you can use the same agent for multiple cloud accounts. For example, you can use the same cloud proxy for a vCenter endpoint or cloud account as for its associated NSX endpoint or cloud account.
- **Communication**: The cloud proxy connects to these domains:
  - Amazon Web Services S3 for cloud proxy OVA download
  - JFrog Artifactory for Docker images
  - Data pipeline service
  - Web API and cloud proxy service
- **Connectivity (ports and protocols)**: The cloud proxy connects to VMware Cloud services through the API gateway or through the data pipeline service. All connections require TLS 1.2 over communications port 443. Cross communication with and between available VMware

Cloud services, such as between VMware Cloud Assembly and vRealize Log Insight Cloud, occurs over HTTPS.

- **Authentication**: User authentication occurs at the VMware Cloud Services Platform (CSP) level, together with Active Directory. The VMware Cloud service validates an incoming authentication header JSON web token (JWT).
- **Credentials and encryption**: A one-time key (OTK) is generated as part of the initial registration workflow for the cloud proxy. The OTK is stored locally in the cloud proxy virtual appliance and expires after 24 hours. A self-signed certificate and an endpoint thumbprint are also stored in the cloud proxy VA.
- **Data**: Encrypted credentials for the VMware Cloud services are stored in the Amazon RDS. Depending on your source VMware Cloud service, the cloud proxy contains these agents:

Agent name	What the agent stores	What the agent collects	Used by	Details
vro-agent	Endpoint credentials.	Workflow and action definitions.	VMware Cloud Assembly	The vRealize Orchestrator agent communicates with the on-premises vRealize Orchestrator server. The vRealize Orchestrator agent propagates information for available workflows to VMware vRealize Automation Cloud and allows triggering of workflow runs from VMware vRealize Automation Cloud. Uses the data pipeline service to communicate with VMware vRealize Automation Cloud. Only outbound traffic from the vRO agent to VMware vRealize Automation Cloud is used.
cloudassembly-cmx-agent	Log files.	Information about PKS/K8S resources.	VMware Cloud Assembly	The CMX agent is responsible for the communication with PKS and K8S clusters.
cloudassembly-blueprint-agent	Temporarily stores request inputs and outputs. Logs are volume mounted.	No automatic collection from on-premises to SaaS. Logs are shared or uploaded only.	VMware Cloud Assembly	Enables VMware Cloud Assembly to integrate with on-premise endpoints such as Ansible and Puppet. Commands are sent over the data pipeline service to communicate with these external accounts.

Agent name	What the agent stores	What the agent collects	Used by	Details
		when manually approved.		
cloudassembly-sddc-agent	Endpoint certificate thumbprint and self-signed certificates.	vCenter and NSX inventory artifacts such as host, machines, storage, networks, and templates.	VMware Cloud Assembly	Information is passed between the VMware Cloud service and the on-premises vCenter Server. The cloud proxy on the vCenter server initiates the connection to the VMware Cloud service. After connected, the cloud proxy receives commands from the VMware Cloud service. A VMware Cloud service cannot initiate the connection to the cloud proxy.
codestream-lemans-agent	Only log files and proxy.properties are volume mounted.	Helps to proxy commands from code stream SaaS to on-premises endpoints.	VMware Code Stream	Enables integration of VMware Code Stream with on-premises endpoints. Helps SaaS execute commands on on-premises endpoints via the data pipeline service. Events such as Gerrit push that happen on on-premises endpoints are pushed to VMware Code Stream to trigger user-configured pipelines.
log-forwarder	Does not store anything.	vCenter and NSX logs.	vRealize Log Insight Cloud	Forwards vCenter and NSX logs to vRealize Log Insight Cloud.

- **Data in cloud proxy**: Self-signed certificates, OTK, private key, public key, and endpoint thumbprints are stored in the cloud proxy VA. Logs generated by the agents are also stored in the cloud proxy.
- **Commands**: The cloud proxy can run these allowed commands:
    - Docker commands

    - Cloud proxy VA upgrade and management trigger scripts such as:
        - /rpm-content/mot.sh
        - /opt/vmware/etc/isv/subsequentboot
        - /rpm-content/subsequentboot.sh
    - VAMI commands
    - Tail command
- **Password expiry**: The root password expires three months after deploying the VA. Upon password expiry, the agent LCM functionality continues to work but the OVA upgrade on reboot cannot occur until the cloud administrator sets a new password.
- **Network proxy**: The network proxy configuration can be provided during cloud proxy VA deployment by using options in the OVA deployment wizard. The cloud administrator can later change the network proxy configuration by using the configure-network-proxy script. The configure-network-proxy script is supplied in the VA's /root folder.

What is the sequential cloud proxy workflow

The cloud proxy installation, deployment, and usage workflow is shown below.

1. Initial user connection to cloud proxy service
    - Description: At initial user connection to the VMware Cloud service, the user login occurs over the VMware Cloud Services Platform (CSP).
    - Authentication and authorization: The username and password are provided to VMWare Identity Manager (vIDM). A JWT token is used for API calls. Permissions are in the JWT.
    - Data in motion: JWT and the session ID.
2. Deploy the cloud proxy OVA, including generation of the one-time-key (OTK)
    - Description: A user connects to the service as part of deploying a cloud proxy. The OTK is generated when a user deploys a cloud proxy for a specific service, for example VMware Cloud Assembly. The cloud proxy service generates the OTK and the call is authenticated against a valid CSP authentication token. The OTK is valid for 24 hours after generation and the cloud proxy must be deployed within this time period.
    - Authentication and authorization: Uses vIDM with the VMware Cloud Services Platform (CSP).
    - Data in motion: JWT and the generated OTK.
3. User download of OVA from VMware Cloud service
    - Description: The OVA is downloaded from an Amazon Web Services S3 bucket. The download URL is provided in the cloud service user interface.
    - Authentication: No authentication or specific authorization is needed to download or deploy the OVA.
    - Data in motion: OVA file.
4. Deployment of the cloud proxy OVA and provisioning process using wizard

- Description: A user deploys the OVA in the target vCenter Server using standard vCenter procedures. During deployment, the cloud proxy generates a self-signed certificate, for example a public/private key pair, and the user is prompted to enter the OTK. The cloud proxy sends the OTK and public key, which is also signed with the private key, and the signature is sent for verification. The service does not verify the signature on this first registration call. Connectivity to the gateway is over TLS 1.2, which is verified by the proxy. The service verifies the OTK by checking the proxy ID and other details before assigning this instance to the respective tenant. The private key is stored in the cloud proxy file system. The cloud proxy service stores the corresponding public key to be used in future signature verification steps.
- Authentication and authorization: OTK, signed message from generated public-private key pair, and TOFU (Trust-On-First-Use).
- Data in motion: OTK, public key of self-signed certificate, signature of request, which is signed by private key, user-provided proxy name, hostname, and IP address.

5. Cloud proxy virtual appliance polls cloud proxy service for commands
   - Description: The cloud proxy polls every 5 seconds for new commands or data. Each request is accompanied by a JWT, which was returned after the user's first registration request, and a request signature. The request signature does not contain the JWT, but the body contains the proxyID. If the JWT is valid, the service returns any pending commands. The proxy stores this JWT in memory. The JWT is valid for 30 minutes.
   - Authentication and authorization: Signature for login and JWT for subsequent calls. If a JWT has expired, the proxy service validates the signature and generates a new JWT token.
   - Data in motion: JWT, proxyID, and request signature signed by private key.
6. Push Docker images for agents
   - Description: When a customer uses a cloud proxy for a new or additional VMware Cloud service, the type of agent is deployed from the VMware Cloud service. This triggers a command push as a response to the proxy that is polling. The command starts a service in the cloud proxy OVA and tells the service the corresponding JFrog details (such as credentials, images to pull, and so on) needed to deploy the agent. The Docker container runs commands and parameters such as mount points are decided from the VMware Cloud service and pushed down the proxy. The proxy runs a sub-service that is responsible for executing these Docker commands and managing the deployed agents.
   - Authentication and authorization: Signature for login and JWT for subsequent calls. If a JWT has expired, the proxy service validates the signature and generates a new JWT token.
   - Data in motion: Commands to execute.
7. Agent connection to the data pipeline service gateway
   - Agents connect to the data pipeline service and send data. Each deployed agent has a built-in data pipeline client that communicates to the data pipeline gateway. This client can either onboard to the data pipeline if it intends to receive commands from the data

pipeline service, or it can simply push data. The data pipeline service client registration process is as follows:

- After you onboard to the service, the cloud proxy service generates an access key for the agent and the customer in the data pipeline service.
- This access key is pushed by the cloud proxy (for example, by using a Docker command line option) when starting the agents. The data pipeline service client in the agent sends the access key to the data pipeline service gateway.
- The access key is verified in the VMware Cloud service and the client is registered.
- As with the cloud proxy workflow, the public key is exchanged and a JWT is returned. The JWT is stored in memory.

  - Authentication and authorization: Signature for login and JWT for subsequent calls. If a JWT has expired, the proxy service validates the signature and generates a new JWT token.
  - Data in motion: Log and metrics data, signature of request, and JWT or access key.

8. Data pipeline service connection to corresponding VMware Cloud service
   - Description: The data pipeline service gateway connects to VMware Cloud services and forwards data to the respective VMware Cloud service. The VMware Cloud service creates streams and routes in the pipeline for agent communication. The configuration information for the streams and routes tells the data pipeline service where to forward the data. The agent is coded to always send its data to the specific stream or route.
   - Authentication and authorization: Signature for login and JWT for subsequent calls. If a JWT has expired, the proxy service validates the signature and generates a new JWT token.
   - Data in motion: Log and metrics data, endpoint-specific information, and data such as vCenter inventory and provisioning requests.
9. Cloud proxy receives commands from VMware Cloud service
   - The cloud proxy receives commands to be executed and credentials to be stored. Below is a list of commands for managing and getting the proxy status remotely from the cloud. These allowed shell scripts are invoked from the cloud and the response is sent to the cloud. As an OS service running as root, the cloud proxy can execute these scripts on the agents.
     - Docker commands
     - Cloud proxy virtual appliance upgrade and management trigger scripts such as:
       - /rpm-content/mot.sh
       - /opt/vmware/etc/isv/subsequentboot
       - /rpm-content/subsequentboot.sh
     - VAMI commands (vCenter Server Appliance Management Interface)
     - Tail command
   - Authentication and authorization: Signature for login and JWT for subsequent calls. If a JWT has expired, the proxy service validates the signature and generates a new JWT token.

- Data in motion: Commands to execute.

10. Cloud proxy service connection to Amazon RDS service
    - Description: VMware Cloud services use an Amazon Web Services Aurora (for example, PostgreSQL) instance. This is an Amazon Web Services service that runs outside the cluster. The schema is flat across tenants. The database credentials are used to communicate through a Hikari connection pool. The credentials are provided to the cloud proxy service by a properties file mounted in a K8S cluster during deployment.
    - Authentication and authorization: Database username and password
    - Data in motion: Commands, proxy and OTK information, and service proxy mappings.
11. Cloud proxy service for agent upgrade commands
    - Description: For agent upgrade, information such as version tag, artifact name, and the new Docker run command is sent as part of the command from the cloud proxy service. JFrog credentials are also pushed. The credentials are read-only. The credentials are used as part of the Docker command and not stored locally on the proxy.
    - Authentication and authorization: Signature for login and JWT for subsequent calls. If a JWT has expired, the proxy service validates the signature and generates a new JWT token.
    - Data in motion: Commands to execute.
12. Agent connection to on-premises devices
    - Description: Service agents connect to on-premises devices and pull data. You enter on-premises device details such as endpoint, username and password, and so on in the VMware Cloud service. These details are stored in the database and sent to the data pipeline service agents. The agents then use these details to call respective APIs.
    - Data in motion: These are specific to the endpoint or cloud account, such as vCenter or NSX-T, in the VMware Cloud service.

General cloud proxy FAQs

What is a cloud proxy?

A cloud proxy connects VMware Cloud services to on-premises networks. It provides multiple capabilities, including a data pipeline service and agent lifecycle management. The cloud proxy is supplied with your VMware Cloud service, for example VMware Cloud Assembly, in the form of a downloadable OVA that you deploy to vCenter Server to create a cloud proxy virtual appliance (VA).

Which VMware Cloud services need a cloud proxy?

The cloud proxy is available for certain types of integrations and endpoints/cloud accounts in VMware Cloud Assembly, VMware Service Broker, VMware Code Stream, and vRealize Log Insight Cloud.

What are some VMware Cloud services scenarios where I need a cloud proxy?

- You are creating or using vCenter, NSX, or VMware Cloud on Amazon Web Services cloud accounts in the VMware Cloud Assembly service.
- You are using any of the VMware Cloud Assembly blueprint components that communicate with on-premises systems, such as on-premises Puppet, vRealize Orchestrator, or Ansible integrations.
- You are integrating with vRealize Orchestrator to run workflows based on event broker subscriptions or as a content source in VMware Service Broker.
- You are provisioning to on-premises PKS (Pivotal Container Services).
- You are using VMware Code Stream and need to integrate with on-premises tools.

What URLs are allowed in the cloud proxy?

- Amazon Web Services S3 to support cloud proxy OVA download.
- JFrog to access Docker images.
- Data pipeline service connection to VMware Cloud services for secure data communication between cloud and on-premises elements.
- Web API and cloud proxy service connection to the VMware Cloud service.

What ports and protocols does the cloud proxy use?

HTTPS over port 443.

Can the cloud proxy be deployed in an HA (high availability) manner?

HA is not available for agents running in the cloud proxy VA. However, you can use the same agent for multiple cloud accounts. For example, you can use the same cloud proxy for a vCenter endpoint or cloud account as for its associated NSX endpoint or cloud account. You can also use vSphere HA for the cloud proxy virtual appliance.

What happens if my cloud proxy can’t connect to the Internet (for example there is an Internet service interruption)?

No data updates are provided to the VMware Cloud service and no commands or actions can be sent from the VMware Cloud service to on-premises endpoints.

Is the data encrypted (at rest and in motion)?

For data in motion, TLS 1.2 or greater encryption is used between the cloud proxy on the on-premises vCenter Server and the endpoint or cloud account in the VMware Cloud service. Communication within the VMware Cloud service is not encrypted. On-premises communication is not encrypted. Data that is at rest in the cloud proxy VA is not encrypted.

Where is data stored in the cloud proxy?

Data that is generated when the cloud proxy VA is started, such as log files, OTK, and private/public certificate pairs, is stored in its file system.

How does the VMware Cloud service know that network traffic is from the cloud proxy? How does it mitigate man-in-the middle attacks?

Public key private key cryptography secures the connection. Trust-on-First-Use (TOFU) is used to mitigate MITM attacks.

How are credentials handled by the VMware Cloud services

See the Data table in the above “How does the cloud proxy support SaaS” section for information about how encrypted credentials for the VMware Cloud service are stored in the Amazon RDS.

VMware vRealize Automation Cloud-specific cloud proxy FAQs

VMware vRealize Automation Cloud is an integrated VMware Cloud services bundle, consisting of VMware Cloud Assembly, VMware Service Broker, and VMware Code Stream.

What data is exchanged between vRealize Automation Cloud and on-premises? What data is kept on vRealize Automation Cloud?

- Request information is passed between the VMware Cloud service and the on-premises vCenter Server.
  The cloud proxy on the vCenter server initiates the connection to the VMware Cloud service. Once connected, the cloud proxy receives commands from the VMware Cloud service. A VMware Cloud service cannot initiate the connection to the cloud proxy.
- vCenter and NSX inventory artifacts such as host, machines, storage, networks, templates and so on are stored in vRealize Automation Cloud.
- Endpoint credentials are stored in vRealize Automation Cloud.
- For details, see the Data table in the above “How does the cloud proxy support SaaS” section.

What system information is stored in vRealize Automation Cloud? Does the cloud proxy VA collect data about networks, machines, and so on, and keep that data locally or is everything relayed to vRealize Automation Cloud?

- Each agent in the cloud proxy VA performs data collection. For example, the cloudassembly-sddc-agent collects inventory artifacts from vCenter. It does not collect data at the OS level.
- vRealize Automation Cloud stores information such as the cloud proxy ID, connection information to the cloud proxy, and customer appliance credentials. The logs contain information such as VM names, IP addresses, and so on.
- Data that is generated during cloud proxy VA startup, such as OTK and private/public certificate pairs, is stored in the file system.

Are any whitelist domains required?

To support the cloud proxy, whitelist access to the following domains is required:

- ci-data-collector.s3.amazonaws.com – enables Amazon Web Services S3 access for cloud proxy OVA download.
- symphony-docker-external.jfrog.io - allows JFrog Artifactory to access Docker images.
- data.mgmt.cloud.vmware.com - enables the data pipeline service connection to VMware Cloud services for secure data communication between cloud and on-premises elements. For non-US

regions, substitute the region value. For example, for the UK, use uk.data.mgmt.cloud.vmware.com and for Japan, use ja.data.mgmt.cloud.vmware.com.

- api.mgmt.cloud.vmware.com – enables the Web API and cloud proxy service connection to VMware Cloud services. For non-US regions, substitute the region value. For example, for the UK, use uk.api.mgmt.cloud.vmware.com and for Japan, use ja.api.mgmt.cloud.vmware.com.
- console.cloud.vmware.com – enables the Web API and cloud proxy service connection to the VMware Cloud service. For non-US regions, substitute the region value. For example, for the UK, use uk.console.cloud.vmware.com.