

Administering vRealize Automation

October 2022

vRealize Automation 8.7

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2022 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 Administering vRealize Automation 5**
- 2 Administering users 6**
 - How do I enable Active Directory groups for projects 7
 - How do I remove users in vRealize Automation 8
 - How do I edit user roles in vRealize Automation 8
 - How do I edit group role assignments in vRealize Automation 9
 - What are the vRealize Automation user roles 10
 - Enable Department of Defense Notice and Consent Banner 25
- 3 Maintaining your appliance 27**
 - Starting and stopping vRealize Automation 27
 - Scale out vRealize Automation from one to three nodes 29
 - Configure an anti-affinity rule and a virtual machine group for a clustered Workspace ONE Access instance 30
 - Replacing an appliance node 31
 - Increase vRealize Automation appliance disk space 32
 - Update the DNS assignment for vRealize Automation 33
 - Change the IP address of a node or cluster 33
 - How do I enable time synchronization 35
 - How do I reset the root password 36
- 4 Using multi-organization tenant configurations in vRealize Automation 38**
 - Set up multi-organization tenancy for vRealize Automation 41
 - Managing certificates and DNS configuration under single-node multi-organization deployments 43
 - Managing certificate and DNS configuration under clustered vRealize Automation deployments 44
 - Logging in to tenants and adding users in vRealize Automation 47
 - Using vRealize Orchestrator with vRealize Automation multi-organization deployments 48
- 5 Working with logs 49**
 - How do I work with logs and log bundles 49
 - How do I configure log forwarding to vRealize Log Insight 53
 - How do I create or update a syslog integration 58
 - How do I delete a syslog integration for logging 59
 - How do I work with content packs 60
- 6 Participating in the Customer Experience Improvement Program 63**

How do I join or leave the program 63

How do I configure the data collection time for the program 64

7 Turning on the in-product feedback form 65

Administering vRealize Automation

1

This guide describes how to monitor and manage critical infrastructure and user management aspects of a vRealize Automation deployment.

The tasks described herein are vital to keeping a vRealize Automation deployment operating appropriately. These tasks include user and group management, and monitoring system logs.

In addition, it describes how to configure and manage multi-organization deployments.

While some vRealize Automation administration tasks are completed from within vRealize Automation, others require the use of related products such as vRealize Suite Lifecycle Manager and Workspace ONE Access. Users should familiarize themselves with these products and their functionality before completing applicable tasks.

For example, for information about backup, restore, and disaster recovery, see the **Backup and Restore, and Disaster Recovery > 2019** section of [vRealize Suite product documentation](#).

Note Disaster recovery is supported in vRealize Automation 8.0.1 and later.

For information about working with vRealize Suite Lifecycle Manager installation, upgrade, and management, see [Lifecycle Manager product documentation](#).

Administering Users and Groups in vRealize Automation

2

vRealize Automation uses VMware Workspace ONE Access, the VMware supplied identity management application to import and manage users and groups. After users and groups are imported or created, you can manage the role assignments for single tenant deployments using the Identity & Access Management page.

vRealize Automation is installed using VMware Lifecycle Manager (vRSLCM or LCM). When installing vRealize Automation you must import an existing Workspace ONE Access instance, or deploy a new one to support identity management. These two scenarios define your management options.

- If you deploy a new Workspace ONE Access instance, you can manage users and groups via LCM. During installation, you can set up an Active Directory connection using Workspace ONE Access. Alternatively, you can view and edit some aspects of users and groups within vRealize Automation using the Identity & Access Management page as described herein.
- If you use an existing Workspace ONE Access instance, you import it for use with vRealize Automation via LCM during installation. In this case, you can continue to use Workspace ONE Access to manage users and groups, or you can use the management functions in LCM.

See [Logging in to tenants and adding users in vRealize Automation](#) for more information about managing users under a multi-organization deployment.

vRealize Automation users must be assigned roles. Roles define access to features within the application. When vRealize Automation is installed with a Workspace ONE Access instance, a default organization is created and the installer is assigned the Organization Owner role. All other vRealize Automation roles are assigned by the Organization Owner.

There are three types of roles in vRealize Automation: organization roles, service roles, and project roles. For Cloud Assembly, Service Broker and Code Stream, typically, user level roles can use resources, while admin level roles are required to create and configure resources. Organizational roles define permissions within the tenant; organizational owners have admin level permissions while organizational members have user level permissions. Organization owners can add and manage other users.

| Organization Roles | Service Roles |
|---|---|
| <ul style="list-style-type: none"> ■ Organization Owner ■ Organization Member | <ul style="list-style-type: none"> ■ Cloud Assembly Administrator ■ Cloud Assembly User ■ Cloud Assembly Viewer ■ Service Broker Administrator ■ Service Broker User ■ Service Broker Viewer ■ Code Stream Administrator ■ Code Stream User ■ Code Stream Viewer |

In addition, there are two main project level roles not shown in the table: Project Administrator, and Project User. These roles are assigned ad hoc on a per project basis with Cloud Assembly. These roles are somewhat fluid. The same user can be an administrator on one project and a user on another project. For more information, see [What are the vRealize Automation user roles](#).

For more information about working with vRealize Suite Lifecycle Manager and Workspace ONE Access, see the following.

This chapter includes the following topics:

- [How do I enable Active Directory groups in vRealize Automation for projects](#)
- [How do I remove users in vRealize Automation](#)
- [How do I edit user roles in vRealize Automation](#)
- [How do I edit group role assignments in vRealize Automation](#)
- [What are the vRealize Automation user roles](#)
- [Enable Department of Defense Notice and Consent Banner](#)

How do I enable Active Directory groups in vRealize Automation for projects

If a group is not available on the Add Groups page when you are adding users to projects, check the Identity & Access Management page and add the group if it is available. If the group is not listed on the Identity & Access Management page in vRealize Automation, the group may not be synchronized in your Workspace One Access instance. You can verify that it has been synchronized and then use this procedure to add the group as shown herein.

To add members of an Active Directory group to a project, you must ensure that the group is synchronized with your Workspace One Access instance and that the group is added to the organization.

Prerequisites

If the groups are not synchronized, they are not available when you try to add them to a project. Verify that you synchronized your Active Directory groups with your Lifecycle Manager instance.

Procedure

- 1 Log in to vRealize Automation as a user from the same Active Directory domain that you are adding. For example, @mycompany.com
- 2 In Cloud Assembly, click Identity & Access Management in the header right navigation.
- 3 Click **Enterprise Groups**, and then click **Assign Roles**.
- 4 Use the search function to find the group that you are adding and select it.
- 5 Assign an organization role.

At a minimum, the group must have an Organization Member role. See [What are the vRealize Automation Cloud Assembly user roles](#) for more information.

- 6 Click **Add Service Access**, add one or more services, and select a role for each.
- 7 Click **Assign**.

Results

You can now add the Active Directory group to a project.

How do I remove users in vRealize Automation

You can remove users as needed in vRealize Automation.

All users are listed by default and you cannot add users with the Identity and Access Management page. You can delete users.

Procedure

- 1 Select the Active Users tab on the Identity & Access Management page.
- 2 Locate and select the users that you want to delete.
- 3 Click **Remove Users**.

Results

The selected users are removed.

How do I edit user roles in vRealize Automation

You can edit roles assigned to Workspace One Access users that have been imported into vRealize Automation.

Prerequisites

Procedure

- 1 In Cloud Assembly, click Identity & Access Management in the header right navigation.
- 2 Select the desired user on the Active Users tab and click **Edit Roles**.
- 3 You can edit the organization and service roles for the user.
 - Select the drop down beside the Assign Organization Roles heading to change the user's relationship with the organization.
 - Click Add Service Access to add new service roles for the user.
 - To remove user roles, click the X beside the applicable service.
- 4 Click **Save**.

Results

The user role assignment is updated as specified.

How do I edit group role assignments in vRealize Automation

You can edit role assignments for groups in vRealize Automation

Prerequisites

Users and groups have been imported from a valid vIDM instance that is associated with your vRealize Automation deployment.

Procedure

- 1 In Cloud Assembly, click Identity & Access Management in the header right navigation.
- 2 Select the Enterprise Groups tab.
- 3 Type the name of the group for which you want to edit role assignments in the search field.
- 4 Edit the role assignments for the selected group. You have two options.
 - Assign Organization Roles
 - Assign Service Roles
- 5 Click **Assign**.

Results

Role assignments are updated as specified.

What are the vRealize Automation user roles

As a organization owner, you can assign users organization roles and service roles. The roles determine what the users can do or see. Then, in the services, the service administrator can assign project roles. To determine the role that you want to assign, evaluate the tasks in the following tables.

Cloud Assembly Service Roles

The Cloud Assembly service roles determine what you can see and do in Cloud Assembly. These service roles are defined in the console by an organization owner.

Table 2-1. Cloud Assembly Service Role Descriptions

| Role | Description |
|------------------------------|---|
| Cloud Assembly Administrator | A user who has read and write access to the entire user interface and API resources. This is the only user role that can see and do everything, including add cloud accounts, create new projects, and assign a project administrator. |
| Cloud Assembly User | A user who does not have the Cloud Assembly Administrator role. In a Cloud Assembly project, the administrator adds users to projects as project members, administrators, or viewers. The administrator can also add a project administrator. |
| Cloud Assembly Viewer | A user who has read access to see information but cannot create, update, or delete values. This is a read-only role across all projects. Users with the viewer role can see all the information that is available to the administrator. They cannot take any action unless you make them a project administrator or a project member. If the user is affiliated with a project, they have the permissions related to the role. The project viewer would not extend their permissions the way that the administrator or member role does. |

In addition to the service roles, Cloud Assembly has project roles. Any project is available in all of the services.

The project roles are defined in Cloud Assembly and can vary between projects.

In the following tables, which tells you what the different service and project roles can see and do, remember that the service administrators have full permission on all areas of the user interface.

The descriptions of project roles will help you decide what permissions to give your users.

- Project administrators leverage the infrastructure that is created by the service administrator to ensure that their project members have the resources they need for their development work.

- Project members work within their projects to design and deploy cloud templates. In the following table, your projects can include only resources that you own or resources that are shared with other project members.
- Project viewers are restricted to read-only access, except in a few cases where they can do non-destructive things like download cloud templates.
- Project supervisors are approvers in Service Broker for their projects where an approval policy is defined with a project supervisor approver. To provide the supervisor with context for approvals, consider also granting them the project member or viewer role.

Table 2-2. Cloud Assembly service roles and project roles

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|------------------------------|---|------------------------------|-----------------------|--|--------------------|--------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Access Cloud Assembly | | | | | | | |
| Console | In the vRA console, you can see and open Cloud Assembly | Yes | Yes | Yes | Yes | Yes | Yes |
| Infrastructure | | | | | | | |
| | See and open the Infrastructure tab | Yes | Yes | Yes | Yes | Yes | Yes |
| Configure - Projects | Create projects | Yes | | | | | |
| | Update, or delete values from project summary, provisioning, Kubernetes, integrations, and test project configurations. | Yes | | | | | |
| | Add users and groups, and assign roles in projects. | Yes | | Yes. Your projects. | | | |
| | View projects | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | Yes. Your projects |
| Configure - Cloud Zones | Create, update, or delete cloud zones | Yes | | | | | |
| | View cloud zones | Yes | Yes | | | | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|------------------------------|--|------------------------------|-----------------------|--|----------------|----------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| | View cloud zone Insights dashboard | Yes | Yes | | | | |
| | View cloud zones alerts | Yes | Yes | | | | |
| Configure - Kubernetes Zones | Create, update, or delete Kubernetes zones | Yes | | | | | |
| | View Kubernetes zones | Yes | Yes | | | | |
| Configure - Flavors | Create, update, or delete flavors | Yes | | | | | |
| | View flavors | Yes | Yes | | | | |
| Configure - Image Mappings | Create, update, or delete image mappings | Yes | | | | | |
| | View image mappings | Yes | Yes | | | | |
| Configure - Network Profiles | Create, update, or delete network profiles | Yes | | | | | |
| | View image network profiles | Yes | Yes | | | | |
| Configure - Storage Profiles | Create, update, or delete storage profiles | Yes | | | | | |
| | View image storage profiles | Yes | Yes | | | | |
| Configure - Pricing Cards | Create, update, or delete pricing cards | Yes | | | | | |
| | View the pricing cards | Yes | Yes | | | | |
| Configure - Tags | Create, update, or delete tags | Yes | | | | | |
| | View tags | Yes | Yes | | | | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|------------------------|---|------------------------------|-----------------------|--|--------------------|--------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Resources - Compute | Add tags to discovered compute resources | Yes | | | | | |
| | View discovered compute resources | Yes | Yes | | | | |
| Resources - Networks | Modify network tags, IP ranges, IP addresses | Yes | | | | | |
| | View discovered network resources | Yes | Yes | | | | |
| Resources - Security | Add tags to discovered security groups | Yes | | | | | |
| | View discovered security groups | Yes | Yes | | | | |
| Resources - Storage | Add tags to discovered storage | Yes | | | | | |
| | View storage | Yes | Yes | | | | |
| Resources - Kubernetes | Deploy or add Kubernetes clusters, and create or add namespaces | Yes | | | | | |
| | View Kubernetes clusters and namespaces | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| Activity - Requests | Delete deployment request records | Yes | | | | | |
| | View deployment request records | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| Activity - Event Logs | View event logs | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|------------------------------|---|------------------------------|-----------------------|--|----------------|--------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Connections - Cloud Accounts | Create, update, or delete cloud accounts | Yes | | | | | |
| | View cloud accounts | Yes | Yes | | | | |
| Connections - Integrations | Create, update, or delete integrations | Yes | | | | | |
| | View integrations | Yes | Yes | | | | |
| Onboarding | Create, update, or delete onboarding plans | Yes | | | | | |
| | View onboarding plans | Yes | Yes | | | Yes. Your projects | |
| Extensibility | | | | | | | |
| | See and open the Extensibility tab | Yes | Yes | | | Yes | |
| Events | View extensibility events | Yes | Yes | | | | |
| Subscriptions | Create, update, or delete extensibility subscriptions | Yes | | | | | |
| | Deactivate subscriptions | Yes | | | | | |
| | View subscriptions | Yes | Yes | | | | |
| Library - Event topics | View event topics | Yes | Yes | | | | |
| Library - Actions | Create, update, or delete extensibility actions | Yes | | | | | |
| | View extensibility actions | Yes | Yes | | | | |
| Library - Workflows | View extensibility workflows | Yes | Yes | | | | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|--------------------------|--|------------------------------|-----------------------|--|-----------------------|-----------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Activity - Action Runs | Cancel or delete extensibility action runs | Yes | | | | | |
| | View extensibility action runs | Yes | Yes | | | Yes. Your projects | |
| Activity - Workflow Runs | View extensibility workflow runs | Yes | Yes | | | | |
| Design | | | | | | | |
| Design | Open the Design tab | Yes | Yes | Yes. | Yes. | Yes. | Yes |
| Cloud Templates | Create, update, and delete cloud templates | Yes | | Yes. Your projects | Yes. Your projects | | |
| | View cloud templates | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| | Download cloud templates | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| | Upload cloud templates | Yes | | Yes. Your projects | Yes. Your projects | | |
| | Deploy cloud templates | Yes | | Yes. Your projects | Yes. Your projects | | |
| | Version and restore cloud templates | Yes | | Yes. Your projects | Yes. Your projects | | |
| | Release cloud templates to the catalog | Yes | | Yes. Your projects | Yes. Your projects | | |
| Custom Resources | Create, update or delete custom resources | Yes | | | | | |
| | View custom resources | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|---------------------------|--|------------------------------|-----------------------|--|--------------------|--------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Custom Actions | Create, update, or delete custom actions | Yes | | | | | |
| | View custom actions | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| Resources | | | | | | | |
| | See and open the Resources tab | Yes | Yes | Yes | Yes | Yes | Yes |
| Deployments | View deployments including deployment details, deployment history, price, monitor, alerts, optimize, and troubleshooting information | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| | Manage alerts | Yes | | Yes. Your projects | Yes. your projects | | |
| | Run day 2 actions on deployments based on policies | Yes | | Yes. Your projects | Yes. Your projects | | |
| Resources - All Resources | View all discovered resources | Yes | Yes | | | | |
| | Run day 2 actions on discovered resources. Actions available only on machines and limited to power on and off for all machines, and remote console for vSphere machines. | Yes | | | | | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|------------------------------|---|------------------------------|-----------------------|--|---------------------|---------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Resources - All Resources | View deployed, onboarded, migrated resources | Yes | Yes | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |
| | Run Day 2 actions on deployed, onboarded, and migrated resources based on policies | Yes | Yes | Yes. Your projects. | Yes. Your projects. | | |
| Resources - Virtual Machines | View discovered machines | Yes | Yes | | | | |
| | Run day 2 actions on discovered machines. Actions are limited to power on and off, and remote console for vSphere machines. | Yes | | | | | |
| | Create New VM | Yes | | | | | |
| | View deployed, onboarded, and migrated resources. | Yes | | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |
| | Run day 2 actions on deployed, onboarded, and migrated resources based on policies | Yes | | Yes. Your projects. | Yes. Your projects. | | |
| Resources - Volumes | View discovered volumes | Yes | Yes | | | | |
| | No day 2 actions available | | | | | | |
| | View deployed, onboarded, and migrated volumes | Yes | Yes | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |

Table 2-2. Cloud Assembly service roles and project roles (continued)

| UI Context | Task | Cloud Assembly Administrator | Cloud Assembly Viewer | Cloud Assembly User User must be a project administrator or member to see and do project-related tasks. | | | |
|------------------------------------|--|------------------------------|-----------------------|--|---------------------|---------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| | Run day 2 actions on deployed, onboarded, and migrated volumes based on policies | Yes | | Yes. Your projects. | Yes. Your projects. | | |
| Resources - Networkin and Security | View discovered networks, load balancers, and security groups | Yes | Yes | | | | |
| | No day 2 actions available | | | | | | |
| | View deployed, onboarded, and migrated networks, load balancers, and security groups | Yes | Yes | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |
| | Run day 2 actions on deployed, onboarded, and migrated networks, load balancers, and security groups based on policies | Yes | | Yes. Your projects. | Yes. Your projects. | | |
| Alerts | | | | | | | |
| | See and open the Alerts tab | Yes | Yes | Yes | Yes | Yes | |
| | Manage alerts | Yes | | Yes. Your projects | Yes. Your projects | | |
| | View alerts | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |

Service Broker Service Roles

The Service Broker service roles determine what you can see and do in Service Broker. These service roles are defined in the console by an organization owner.

Table 2-3. Service Broker Service Role Descriptions

| Role | Description |
|------------------------------|---|
| Service Broker Administrator | Must have read and write access to the entire user interface and API resources. This is the only user role that can perform all tasks, including creating a new project and assigning a project administrator. |
| Service Broker User | Any user who does not have the Service Broker Administrator role. In a Service Broker project, the administrator adds users to projects as project members, administrators, or viewers. The administrator can also add a project administrator. |
| Service Broker Viewer | A user who has read access to see information but cannot create, update, or delete values. Users with the viewer role can see all the information that is available to the administrator. They cannot take any action unless you make them a project administrator or a project member. If the user is affiliated with a project, they have the permissions related to the role. The project viewer would not extend their permissions the way that the administrator or member role does. |

In addition to the service roles, Service Broker has project roles. Any project is available in all of the services.

The project roles are defined in Service Broker and can vary between projects.

In the following tables, which tells you what the different service and project roles can see and do, remember that the service administrators have full permission on all areas of the user interface.

Use the following descriptions of project roles will help you as you decide what permissions to give your users.

- Project administrators leverage the infrastructure that is created by the service administrator to ensure that their project members have the resources they need for their development work.
- Project members work within their projects to design and deploy cloud templates. In the following table, Your projects can include only resources that you own or resources that are shared with other project members.
- Project viewers are restricted to read-only access.
- Project supervisors are approvers in Service Broker for their projects where an approval policy is defined with a project supervisor approver. To provide the supervisor with context for approvals, consider also granting them the project member or viewer role.

Table 2-4. Service Broker Service Roles and Project Roles

| UI Context | Task | Service Broker Administrator | Service Broker Viewer | Service Broker User User must be a project administrator to see and do project-related tasks. | | | |
|------------------------------|---|------------------------------|-----------------------|--|--------------------|--------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Access Service Broker | | | | | | | |
| Console | In the console, you can see and open Service Broker | Yes | Yes | Yes | Yes | Yes | Yes |
| Infrastructure | | | | | | | |
| | See and open the Infrastructure tab | Yes | Yes | | | | |
| Configure - Projects | Create projects | Yes | | | | | |
| | Update, or delete values from project summary, provisioning, Kubernetes, integrations, and test project configurations. | Yes | | | | | |
| | Add users and groups, and assign roles in projects. | Yes | | Yes. Your projects. | | | |
| | View projects | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| Configure - Cloud Zones | Create, update, or delete cloud zones | Yes | | | | | |
| | View cloud zones | Yes | Yes | | | | |
| Configure - Kubernetes Zones | Create, update, or delete Kubernetes zones | Yes | | | | | |
| | View Kubernetes zones | Yes | Yes | | | | |
| Connections - Cloud Accounts | Create, update, or delete cloud accounts | Yes | | | | | |

Table 2-4. Service Broker Service Roles and Project Roles (continued)

| UI Context | Task | Service Broker Administrator | Service Broker Viewer | Service Broker User User must be a project administrator to see and do project-related tasks. | | | |
|-----------------------------|--|------------------------------|-----------------------|--|----------------|----------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| | View cloud accounts | Yes | Yes | | | | |
| Connections - Integrations | Create, update, or delete integrations | Yes | | | | | |
| | View integrations | Yes | Yes | | | | |
| Activity - Requests | Delete deployment request records | Yes | | | | | |
| | View deployment request records | Yes | | | | | |
| Activity - Event Logs | View event logs | Yes | | | | | |
| Content and Policies | | | | | | | |
| | See and open the Content and Policies tab | Yes | Yes | | | | |
| Content Sources | Create, update, or delete content sources | Yes | | | | | |
| | View content sources | Yes | Yes | | | | |
| Content Sharing | Add or remove shared content | Yes | | | | | |
| | View shared content | Yes | Yes | | | | |
| Content | Customize form and configure item | Yes | | | | | |
| | View content | Yes | Yes | | | | |
| Policies - Definitions | Create, update, or delete policy definitions | Yes | | | | | |
| | View policy definitions | Yes | Yes | | | | |
| Policies - Enforcement | View enforcement log | Yes | Yes | | | | |

Table 2-4. Service Broker Service Roles and Project Roles (continued)

| UI Context | Task | Service Broker Administrator | Service Broker Viewer | Service Broker User User must be a project administrator to see and do project-related tasks. | | | |
|------------------------------|---|------------------------------|-----------------------|--|--------------------|--------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| Notifications - Email Server | Configure an email server | Yes | | | | | |
| Catalog | | | | | | | |
| | See and open the Catalog tab | Yes | Yes | Yes | Yes | Yes | Yes |
| | View available catalog items | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| | Request a catalog item | Yes | | Yes. Your projects | Yes. Your projects | | |
| Resources | | | | | | | |
| | See and open the Resources tab | Yes | Yes | Yes. | Yes | Yes | Yes |
| Deployments | View deployments, including deployment details, deployment history, price, monitor, alerts, optimize, and troubleshooting information | Yes | Yes | Yes. Your projects | Yes. Your projects | Yes. Your projects | |
| | Manage alerts | Yes | | Yes. Your projects | Yes. Your projects | | |
| | Run day 2 actions on deployments based on policies | Yes | | Yes. Your projects | Yes. Your projects | | |
| Resources - All Resources | View all discovered resources | Yes | Yes | | | | |

Table 2-4. Service Broker Service Roles and Project Roles (continued)

| UI Context | Task | Service Broker Administrator | Service Broker Viewer | Service Broker User | | | |
|------------------------------|---|------------------------------|-----------------------|---|---------------------|---------------------|--------------------|
| | | | | User must be a project administrator to see and do project-related tasks. | | | |
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| | Run day 2 actions on discovered resources. Actions available only on machines and limited to power on and off for all machines, and remote console for vSphere machines. | Yes | | | | | |
| Resources - All Resources | View deployed, onboarded, migrated resources | Yes | Yes | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |
| | Run Day 2 actions on deployed, onboarded, and migrated resources based on policies | Yes | Yes | Yes. Your projects. | Yes. Your projects. | | |
| Resources - Virtual Machines | View discovered machines | Yes | Yes | | | | |
| | Run day 2 actions on discovered machines. Actions are limited to power on and off, and remote console for vSphere machines. | Yes | | | | | |
| | Create New VM | Yes | | | | | |
| | View deployed, onboarded, and migrated resources. | Yes | | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |

Table 2-4. Service Broker Service Roles and Project Roles (continued)

| UI Context | Task | Service Broker Administrator | Service Broker Viewer | Service Broker User User must be a project administrator to see and do project-related tasks. | | | |
|-------------------------------------|--|------------------------------|-----------------------|--|---------------------|---------------------|--------------------|
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| | Run day 2 actions on deployed, onboarded, and migrated resources based on policies | Yes | | Yes. Your projects. | Yes. Your projects. | | |
| Resources - Volumes | View discovered volumes | Yes | Yes | | | | |
| | No day 2 actions available | | | | | | |
| | View deployed, onboarded, and migrated volumes | Yes | Yes | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |
| | Run day 2 actions on deployed, onboarded, and migrated volumes based on policies | Yes | | Yes. Your projects. | Yes. Your projects. | | |
| Resources - Networking and Security | View discovered networks, load balancers, and security groups | Yes | Yes | | | | |
| | No day 2 actions available | | | | | | |
| | View deployed, onboarded, and migrated networks, load balancers, and security groups | Yes | Yes | Yes. Your projects. | Yes. Your projects. | Yes. Your projects. | |
| | Run day 2 actions on deployed, onboarded, and migrated networks, load balancers, and security groups based on policies | Yes | | Yes. Your projects. | Yes. Your projects. | | |
| Approvals | | | | | | | |

Table 2-4. Service Broker Service Roles and Project Roles (continued)

| UI Context | Task | Service Broker Administrator | Service Broker Viewer | Service Broker User | | | |
|------------|--------------------------------|------------------------------|-----------------------|---|----------------------------------|----------------------------------|--|
| | | | | User must be a project administrator to see and do project-related tasks. | | | |
| | | | | Project Administrator | Project Member | Project Viewer | Project Supervisor |
| | See and open the Approvals tab | Yes | Yes | Yes | Yes | Yes | Yes |
| | Respond to approval requests | Yes | | Yes. Your projects and the policy approver is Project Administrator | Only if you are a named approver | Only if you are a named approver | Yes. Your projects and the policy approver is Project Supervisor |

Enable Department of Defense Notice and Consent Banner

For some government customers an administrator must configure the standard Department of Defense (DoD) notice and consent banner in Workspace ONE Access in order for users to access vRealize Automation.

The Standard Mandatory DoD Notice and Consent Banner text is as follows:

You are accessing a U.S. Government (USG) Information System (IS) that is provided for USG-authorized use only. By using this IS (which includes any device attached to this IS) you consent to the following conditions:

- The USG routinely intercepts and monitors communications on this IS for purposes including, but not limited to, penetration testing, COMSEC monitoring, network operations and defense, personnel misconduct (PM), law enforcement (LE), and counterintelligence (CI) investigations.
- At any time, the USG may inspect and seize data stored on this IS.
- Communications using, or data stored on, this IS are not private, are subject to routine monitoring, interception, and search, and may be disclosed or used for any USG-authorized purpose.

The following steps describe how to configure this banner in Workspace ONE Access. For more information, see the Workspace ONE Access administrative console documentation.

Procedure

- 1 Log in to the Workspace ONE administrative console as an administrator.
- 2 In the VMware Identity Manager Console, click the Identity and Access Management tab.
- 3 Click Setup and then click the Connectors tab.
- 4 Click the Worker link for each connector that you want to configure.

- 5 Click the Auth Adapters tab and then `CertificateAuthAdapter`.
- 6 Click the Enable Consent Form before Authentication check box.
- 7 Paste the Standard Mandatory DoD Notice and Consent Banner text into the Consent Form Content box.
- 8 Save your changes.

Results

Maintaining your vRealize Automation appliance

3

As a system administrator, you might need to perform various tasks to ensure the proper functioning of your installed vRealize Automation application.

If you are just getting started with vRealize Automation, these are not required tasks. Knowing how to perform these tasks is useful if you need to resolve performance or product behavior issues.

This chapter includes the following topics:

- Starting and stopping vRealize Automation
- Scale out vRealize Automation from one to three nodes
- Configure an anti-affinity rule and a virtual machine group for a clustered Workspace ONE Access instance
- Replacing a vRealize Automation appliance node
- Increase vRealize Automation appliance disk space
- Update the DNS assignment for vRealize Automation
- Change IP addresses of vRealize Automation node or cluster
- How do I enable time synchronization of vRealize Automation
- How do I reset the root password for vRealize Automation

Starting and stopping vRealize Automation

Observe the proper procedures when starting or shutting down vRealize Automation.

The recommended manner to shut down and start vRealize Automation components is to use the Power OFF and ON functionality provided in **Lifecycle Operations > Environments** section of vRealize Suite Lifecycle Manager. The following procedures outline manual methods to shut down and start vRealize Automation components in case vRealize Suite Lifecycle Manager is not available for some reason.

Shut down vRealize Automation

To preserve data integrity, shut down the vRealize Automation services before powering off the virtual appliances. Using SSH or VMRC, you can shut down or start all nodes from any individual appliance.

Note Avoid using `vracli reset vidm` commands if at all possible. This command resets all configurations of Workspace One Access and breaks the association between users and provisioned resources.

- 1 Log in to the console of any vRealize Automation appliance using either SSH or VMRC.
- 2 To shut down the vRealize Automation services on all cluster nodes, Run the following set of commands.

Note If you copy any of these commands to run and they fail, paste them into notepad first, and then copy them again before running them. This procedure strips out any hidden characters and other artifacts that might exist in the documentation source.

```
/opt/scripts/deploy.sh --shutdown
```

- 3 Shut down the vRealize Automation appliances.

Your vRealize Automation deployment is now shut down.

Start vRealize Automation

Following an unplanned shutdown, a controlled shutdown, or a recovery procedure, you must restart vRealize Automation components in a specific order. vRLCM is a non-critical component, so you can start it at any time. VMware Workspace ONE Access, formerly VMware Identity Management, components must be started before you start vRealize Automation.

Note Verify that applicable load balancers are running before starting vRealize Automation components.

- 1 Power on all vRealize Automation appliances and wait for them to start.
- 2 Log into the console for any appliance using SSH or VMRC and run the following command to restore the services on all nodes.

```
/opt/scripts/deploy.sh
```

- 3 Verify that all services are up and running with the following command.

```
kubectl get pods --all-namespaces
```

Note You should see three instances of every service, with a status of either Running or Completed.

When all services are listed as Running or Completed, vRealize Automation is ready to use.

Restart vRealize Automation

You can restart all vRealize Automation services centrally from any of the appliances in your cluster. Follow the preceding instructions to shut down vRealize Automation, and then use the instructions to start vRealize Automation. Before restarting vRealize Automation, verify that all applicable load balancer and VMware Workspace ONE Access components are running.

When all services are listed as Running or Completed, then vRealize Automation is ready to use.

Run the following command to verify that all services are running:

```
kubectl -n prelude get pods
```

Scale out vRealize Automation from one to three nodes

As needs expand, you can scale out a vRealize Automation deployment from one node to three nodes.

You must use features of vRealize Suite Lifecycle Manager to complete many steps of this procedure. For information about working with vRealize Suite Lifecycle Manager installation, upgrade, and management, see [Lifecycle Manager product documentation](#).

If you are using a three node clustered deployment, vRealize Automation can typically withstand the failure of one node and still function. The failure of two nodes in a three node cluster will render vRealize Automation non-functional.

Prerequisites

This procedure assumes that you already have a functioning single node vRealize Automation deployment.

Procedure

- 1 Shut down all vRealize Automation appliances.

To shut down the vRealize Automation services on all cluster nodes, run the following set of commands.

```
/opt/scripts/svc-stop.sh  
sleep 120  
/opt/scripts/deploy.sh --onlyClean
```

Now you can shut down the vRealize Automation appliances.

- 2 Take a deployment snapshot.

Use the Create Snapshot option in vRealize Suite Lifecycle Manager **Lifecycle Operations > Environments > vRA > View Details**.

Note Online snapshots, taken without shutting down vRealize Automation nodes are supported from 8.0.1. For vRealize Automation 8.0 environments, you must stop vRealize Automation nodes first.

- 3 Power on the vRealize Automation appliance and bring up all containers.
- 4 Using the Locker functionality located at **LCM > Locker > Certificates** in vRealize Suite Lifecycle Manager, generate or import vRealize Automation certificates for all components including vRealize Suite Lifecycle Manager node FQDNs and the vRealize Automation Load Balancer fully qualified domain name.

Add the names of all three appliances in the Subject Alternative Names.
- 5 Import the new certificate into vRealize Suite Lifecycle Manager.
- 6 Replace the existing vRealize Suite Lifecycle Manager certificate with the one generated in the previous step using the LCM **Lifecycle Operations > Environments > vRA > View Details** Replace Certificate option.
- 7 Scale out vRealize Automation to three nodes using the Add Components selection in **LCM > Lifecycle Operations > Environments > vRA > View Details**.

Results

vRealize Automation has been scaled to a three node deployment.

Configure an anti-affinity rule and a virtual machine group for a clustered Workspace ONE Access instance

If your vRealize Automation environment uses a clustered Workspace ONE Access instance, create an anti-affinity rule and machine cluster to ensure proper vSphere High Availability workflow.

To protect any clustered Workspace ONE Access nodes from a host-level failure, configure an anti-affinity rule to run virtual machines that exist on different hosts in the default vSphere management cluster. After you create an anti-affinity rule, configure a virtual machine group to define the desired machine start-up order. By using a defined machine start-up order, you can ensure that vSphere High Availability powers on the clustered Workspace ONE Access nodes in the correct order for your environment.

For information about how to configure anti-affinity rules and a virtual machine group, see [Configure an Anti-Affinity Rule and a Virtual Machine Group for the Clustered Workspace ONE Access Instance](#) in the [VMware Cloud Foundation Product Documentation](#).

Affinity rule considerations when upgrading from one vRealize Automation release to another

vRealize Suite Lifecycle Manager does not support anti-affinity rules for vRealize Automation 8.x. Because vRealize Suite Lifecycle Manager is used by vRealize Easy Installer during vRealize Automation upgrade, and there is no specific order of powering vRealize Automation nodes off and on during upgrade, issues can occur if the order used conflicts with affinity rules that define the order in which machines are powered off and on. When using vRealize Suite Lifecycle Manager or vRealize Easy Installer to upgrade from one vRealize Automation release to another, disable affinity rules prior to starting the upgrade.

For information about how to upgrade from one vRealize Automation to another, see [Installing vRealize Automation with vRealize Easy Installer](#) in [vRealize Automation Product Documentation](#).

Replacing a vRealize Automation appliance node

When a vRealize Automation appliance in a multiple-node, high availability (HA) configuration has failed, you might need to replace the faulty node.

Caution Before proceeding, VMware recommends that you contact technical support to troubleshoot the HA issue and verify that the problem is isolated to one node.

If technical support determines that you need to replace the node, take the following steps.

- 1 In vCenter, take backup snapshots of every appliance in the HA configuration.
In the backup snapshots, don't include virtual machine memory.
- 2 Shut down the faulty node.
- 3 Make note of the faulty node vRealize Automation software build number, and network settings.

Note the FQDN, IP address, gateway, DNS servers, and especially MAC address. Later, you assign the same values to the replacement node.

- 4 The primary database node must be one of the healthy nodes. Follow these steps:
 - a Log in as root to the command line of a healthy node.
 - b Find the name of the primary database node by running the following command.

```
vracli status | grep primary -B 1
```

The result should be similar to this example, where postgres-1 is the primary database node.

```
"Conninfo":
"host=postgres-1.postgres.prelude.svc.cluster.local
dbname=repmgr-db user=repmgr-db passfile=/scratch/repmgr-db.cred
connect_timeout=10",
"Role": "primary",
```

- c Verify that the primary database node is healthy by running the following command.

```
kubectl -n prelude get pods -o wide | grep postgres
```

The result should be similar to this example, where postgres-1 is in the list as running and healthy.

```
postgres-1 1/1 Running 0 39h 12.123.2.14 vc-vm-224-84.company.com <none> <none>
postgres-2 1/1 Running 0 39h 12.123.1.14 vc-vm-224-85.company.com <none> <none>
```

Important If the primary database node is faulty, contact technical support instead of proceeding.

- 5 From the root command line of the healthy node, remove the faulty node.

```
vracli cluster remove faulty-node-FQDN
```

- 6 Use vCenter to deploy a new, replacement vRealize Automation node.

Deploy the same vRealize Automation software build number, and apply the network settings from the faulty node. Include the FQDN, IP address, gateway, DNS servers, and especially MAC address that you noted earlier.

- 7 Power on the replacement node.
- 8 Log in as root to the command line of the replacement node.
- 9 Verify that the initial boot sequence has finished by running the following command.

```
vracli status first-boot
```

Look for a `First boot complete` message.

- 10 From the replacement node, join the vRealize Automation cluster.

```
vracli cluster join primary-DB-node-FQDN
```

- 11 Log in as root to the command line of the primary database node.
- 12 Deploy the repaired cluster by running the following script.

```
/opt/scripts/deploy.sh
```

Increase vRealize Automation appliance disk space

You might need to increase vRealize Automation appliance disk space for purposes such as log file storage.

Procedure

- 1 Use vSphere to expand the VMDK on the vRealize Automation appliance.
- 2 Log in to the command line of the vRealize Automation appliance as a root user.

- 3 From the command prompt, run the following vRealize Automation command:

```
vracli disk-mgr resize
```

If vRealize Automation resizing fails, see [Knowledge Base article 79925](#).

Update the DNS assignment for vRealize Automation

An administrator can update the DNS assignments for vRealize Automation.

Procedure

- 1 Log in to the console for any vRealize Automation appliance using either SSH or VMRC.
- 2 Run the following command.

```
vracli network dns set --servers DNS1,DNS2
```

- 3 Verify that the new DNS servers were properly applied to all vRealize Automation nodes with `vracli network dns status` command.
- 4 Run the following set of commands to shut down the vRealize Automation services on all cluster nodes, .

```
/opt/scripts/svc-stop.sh
sleep 120
/opt/scripts/deploy.sh --onlyClean
```

- 5 Restart the vRealize Automation nodes and wait for them to start completely.
- 6 Log in to each vRealize Automation node with SSH and verify that the new DNS servers are listed in `/etc/resolv.conf`.
- 7 On one of the vRealize Automation nodes, run the following command to start the vRealize Automation services: `/opt/scripts/deploy.sh`

Results

The vRealize Automation DNS settings are changed as specified.

Change IP addresses of vRealize Automation node or cluster

You can change the IP address of a vRealize Automation node or cluster.

For example, you might want to migrate your deployed vRealize Automation environment to a more convenient vCenter or to support vRealize Automation fail-over.

As a vRealize Automation administrator, you can use the following procedure to set a new IP address for the vRealize Automation node or cluster and then redeploy services to the new IP address.

Note Before you proceed with changing the IP of a vRealize Automation node or cluster, you must verify that the node or cluster is in a healthy state. Attempting to run this procedure on a node or cluster that is not in a healthy state can create problems that are very challenging to resolve.

In this procedure, you will restart vRealize Automation in a specific and sequential manner. For related information about shutting down and restarting vRealize Automation, see [Starting and stopping vRealize Automation](#).

- 1 Verify that the vRealize Automation node or cluster is in a healthy state by using the following command.

```
vracli service status
```

- 2 When vRealize Automation is in healthy state, set the alternative IP of the node or cluster appliance(s) by using the following command.

```
vracli network alternative-ip set --dns DNSIPaddress1,DNSIPaddress2 IPV4_address  
Gateway_IPV4_address
```

If you are working with a cluster, set the alternative IP of each applicable node in the cluster.

- 3 Shut down the services by using the following command.

```
/opt/scripts/deploy.sh -shutdown
```

- 4 If needed, perform a vRealize Automation fail-over or migration operation. See information about [VMware Site Recovery Manager](#) and your own internal procedures and practices.
- 5 Change the IP of vRealize Automation by using the following command.

```
vracli network alternative-ip swap
```

If you are using a vRealize Automation cluster, you must change the IP address of each node in the cluster.

- 6 Reboot vRealize Automation by using the following command.

```
shutdown -r now
```

If you are using a vRealize Automation cluster, you must reboot each node in the cluster.

- 7 Redeploy vRealize Automation services by using the following command.

```
/opt/scripts/deploy.sh
```

After you reboot vRealize Automation and the redeploy services are running, vRealize Automation should be available at the new IP address.

How do I enable time synchronization of vRealize Automation

You can enable time synchronization on your vRealize Automation deployment by using the vRealize Automation Appliance command line.

You can configure time synchronization for your standalone or clustered vRealize Automation deployment by using the Network Time Protocol (NTP) networking protocol. vRealize Automation supports two, mutually exclusive, NTP configurations:

| NTP configuration | Description |
|-------------------|---|
| ESXi | <p>This configuration can be used when the ESXi server hosting the vRealize Automation is synchronized with an NTP server. If you are using a clustered deployment, all ESXi hosts must be synchronized with an NTP server. For more information about configuring NTP for ESXi, see KB article 57147 Configuring Network Time Protocol (NTP) on an ESXi host using the vSphere Web Client.</p> <hr/> <p>Note If your vRealize Automation deployment is migrated to a ESXi host that is not synchronized to an NTP server, you can experience clock drift.</p> |
| systemd | <p>This configuration uses the systemd-timesyncd daemon to synchronize the clocks of your vRealize Automation deployment.</p> <hr/> <p>Note By default, the systemd-timesyncd daemon is enabled, but configured with no NTP servers. If the vRealize Automation Appliance uses a dynamic IP configuration, the appliance can use any NTP servers received by the DHCP protocol.</p> |

Procedure

- 1 Log in to the vRealize Automation Appliance command line as **root**.
- 2 Enable NTP with ESXi.
 - a Run the `vracli ntp esxi` command.
 - b (Optional) To confirm the status of the NTP configuration, run the `vracli ntp status` command.

You can also reset the NTP configuration to the default state by running the `vracli ntp reset` command.

3 Enable NTP with systemd.

- a Run the `vracli ntp systemd --set FQDN_or_IP_of_systemd_server` command.

Note You can add multiple systemd NTP servers by separating their network addresses with a comma. Each network address must be placed inside single quotation marks. For example, `vracli ntp systemd --set 'ntp_address_1','ntp_address_2'`

- b (Optional) To confirm the status of the NTP configuration, run the `vracli ntp status` command.

Results

You have enabled time synchronization for your vRealize Automation Appliance deployment.

What to do next

The NTP configuration can fail if there is a time difference of above 10 minutes between the NTP server and the vRealize Automation deployment. To resolve this problem, reboot the vRealize Automation Appliance.

How do I reset the root password for vRealize Automation

You can reset a lost or forgotten vRealize Automation root password.

In this procedure, you use a command line window on the host vCenter appliance to reset your organization's vRealize Automation root password.

Prerequisites

This process is for vRealize Automation administrators and requires the credentials needed to access the host vCenter appliance.

Procedure

- 1 Shut down and start up vRealize Automation by using the procedure described in [Starting and stopping vRealize Automation](#).
- 2 When the Photon operating system command line window appears, enter `e` and press the **Enter** key to open the GNU GRUB boot menu editor.

- 3 In the GNU GRUB editor, enter `rw init=/bin/bash` at the end of the line that begins with `linux` `"/" $photon_linux root=rootpartition` as shown below:

```

vcenter
GNU GRUB version 2.02~beta2

setparams 'Photon'

linux "/"$photon_linux root=$rootpartition not ifnames=0 $photon_cmd\
line coredump_filter=0x37 consoleblank=0 rw init=/bin/bash_
if [ "$photon_initrd" ]; then
  initrd "/"$photon_initrd
fi

Minimum Emacs-like screen editing is supported. TAB lists
completions. Press Ctrl-x or F10 to boot, Ctrl-c or F2 for a
command-line or ESC to discard edits and return to the GRUB
menu.

```

- 4 Click the **F10** key to push your change and restart vRealize Automation.
- 5 Wait for vRealize Automation to restart.
- 6 At the `root [/]#` prompt, enter `passwd` and press the **Enter** key.
- 7 At the `New password:` prompt, enter your new password and press the **Enter** key.
- 8 At the `Retype new password:` prompt, reenter your new password and press the **Enter** key.
- 9 At the `root [/]#` prompt, enter `reboot -f` and press the **Enter** key to complete the root password reset process.

```

root [ / ]# passwd
New password:
Retype new password:
passwd: password updated successfully
root [ / ]# reboot -f_

```

What to do next

As a vRealize Automation administrator, you can now log in to vRealize Automation with the new root password.

Using multi-organization tenant configurations in vRealize Automation

4

vRealize Automation enables IT providers, to set up multiple tenants, or organizations, within each deployment. Providers can set up multiple tenant organizations and allocate infrastructure within each deployment. Providers can also manage users for tenants. Each tenant manages its own projects, resources, and deployments.

In a vRealize Automation multi-organization configuration, providers can create multiple organizations, and each tenant organization uses its own projects, resources and deployments. While providers cannot manage tenant infrastructure remotely, they can log in to tenants and manage infrastructure within their tenants.

Multi-tenancy relies on coordination and configuration of three different VMware products as outlined below:

- Workspace ONE Access - This product provides the infrastructure support for multi-tenancy and the Active Directory domain connections that provide user and group management within tenant organizations.
- vRealize Suite Lifecycle Manager - This product supports the creation and configuration of tenants for supported products, such as vRealize Automation. In addition, it provides some certificate management capabilities.
- vRealize Automation - Providers and users log in to vRealize Automation to access tenants in which they create and manage deployments.

When configuring multi-tenancy, users should be familiar with all three of these products and their associated documentation.

For more information about working with vRealize Suite Lifecycle Manager and Workspace ONE Access, see the following.

- vRealize Suite Lifecycle Manager- see [Lifecycle Manager product documentation](#)
- Workspace ONE Access - see [User Management with VMware Identity Manager and VMware Workspace ONE Access Administration](#)

Administrators with vRealize Suite Lifecycle Manager privileges create and manage tenants using the Lifecycle Manager Tenants page located under the Identity and Tenant Management service. Tenants are constructed using an Active Directory IWA or LDAP connection, and they are supported by the associated VMware Workspace ONE Access instance that is required for vRealize Automation deployments. See the associated documentation for information about using Lifecycle Manager.

When configuring multi-tenancy, you start with a base, or master tenant. This tenant is the default tenant that is created when the underlying Workspace ONE Access application is deployed. Other tenants, known as sub-tenants, can be based upon the master tenant. vRealize Automation currently supports up to 20 tenant organizations with the standard three node deployment.

Before enabling vRealize Automation for multi-tenancy, you must first install the application in a single organization configuration, and then use Lifecycle Manager to set up a multi-organization configuration. A Workspace ONE Access deployment supports the management of tenants and the associated Active Directory domain connections.

When you initially set up multi-tenancy, a provider administrator is designated in Lifecycle Manager. You can change this designation or add administrators later if desired. Under multi-organization configurations, vRealize Automation users and groups are managed primarily through Workspace ONE Access.

After organizations are created, authorized users can log in to their applications to create or work with projects and resources and create deployments. Administrators can manage user roles in vRealize Automation.

Setting up for a multi-organization configuration

You can enable a multi-organization deployment after completing a vRealize Automation installation. When setting up a multi-organization configuration, you must configure your external Workspace ONE Access for multi-tenancy use and then use Lifecycle manager to create and configure tenants. This applies to both new and existing deployments. As an initial step to setting up tenants, you must use Lifecycle Manager to set an alias for the master tenant that was created by default on Workspace ONE Access. Sub-tenants that you create based on this master tenant inherit the Active Directory domain configurations from this master tenant.

In Lifecycle Manager, you assign tenants to a product, such as vRealize Automation, and to a specific environment. When setting up a tenant, you must also designate a tenant administrator. By default, multi-tenancy is enabled based on tenant hostname. Users can elect to manually configure tenant name by DNS name. During this procedure you must set several flags to support multi-tenancy, and you must configure the load balancer as well.

If you use a clustered instance, both the Workspace ONE Access and vRealize Automation tenant based hostnames will point to the load balancer.

If your clustered vRealize Automation and Workspace ONE Access load balancers do not use wildcard certificates, then users must add tenant hostnames as SAN entries on the certificates. for every new tenant that is created.

You cannot delete tenants in vRealize Automation or in Lifecycle Manager. If you need to add tenants to an existing multi-tenancy deployment, you can do this using Lifecycle Manager, but it will necessitate downtime of three to four hours.

Refer to the documentation links at the beginning of this topic for more information about using vRealize Suite Lifecycle Manager Workspace ONE Access.

Hostnames and multi-tenancy

In prior versions of vRealize Automation, users accessed tenants with URLs that were based on directory path. In the current multi-tenancy implementation, users access tenants based on hostname.

Also, the hostname format that vRealize Automation users will use to access tenants differs from the format that is used to access tenants within Workspace ONE Access. For example, a valid hostname would look like the following: `tenant1.example.eng.vmware.com` as opposed to `vidm-node1.eng.vmware.com`.

Multi-tenancy and certificates

You must create certificates for all components involved in a multi-organization configuration. You will need one or more certificates for Workspace ONE Access, Lifecycle Manager, and vRealize Automation, depending on whether you are using a single node configuration or a clustered configuration.

When configuring certificates, you can use either wildcards with the SAN names or dedicated names. Using wildcards will simplify certificate management somewhat as certificates must be updated whenever you add new tenants. If your vRealize Automation and Workspace ONE Access load balancer do not use wildcard certificates, then you must add tenant hostnames as SAN entries on the certificates for every new tenant that is created. Also, if you use SAN, certificates must be updated manually if you add or delete hosts or change a hostname. You must also update DNS entries for tenants.

Note that Lifecycle Manager does not create separate certificates for each tenant. Instead it creates a single certificate with each tenant hostname listed. For basic configurations, the tenant's CNAME uses the following format: `tenantname.vrahostname.domain`. For high availability configurations, the name uses the following format: `tenantname.vraLBhostname.domain`.

If you are using a clustered Workspace ONE Access configuration, note that Lifecycle Manager cannot update the load balancer certificate, so you must update it manually. Also, if you need to re-register products or services that are external to Lifecycle Manager, this is a manual process.

This chapter includes the following topics:

- [Set up multi-organization tenancy for vRealize Automation](#)
- [Logging in to tenants and adding users in vRealize Automation](#)
- [Using vRealize Orchestrator with vRealize Automation multi-organization deployments](#)

Set up multi-organization tenancy for vRealize Automation

You can set up multi-organization tenancy for vRealize Automation using vRealize Suite Lifecycle Manager.

The following is a high level description of the procedure to set up multi-tenancy for vRealize Automation including configuring DNS and certificates. It focuses on a single node deployment but includes notes for a clustered configuration.

See <https://vmwarelab.org/2020/04/14/vrealize-automation-8-1-multi-tenancy-setup-with-vrealize-suite-lifecycle-manager-8-1/> for more information and a video demonstration of configuring a vRealize Automation multi-organization configuration.

Prerequisites

- Install and configure Workspace ONE Access version 3.3.4 or newer.
- Install and configure vRealize Suite Lifecycle Manager version 8.5.

Procedure

- 1 Create the required A and CNAME Type DNS records.
 - For your master tenant and each sub-tenant, you must create and apply a SAN certificate.
 - For single node deployments, the vRealize Automation FQDN points to the vRealize Automation appliance, and the Workspace ONE Access FQDN points to the Workspace ONE Access appliance.
 - For clustered deployments, both the Workspace ONE Access and vRealize Automation tenant-based FQDNs must point to their respective load balancers. Workspace ONE Access is configured with SSL Termination, so the certificate is applied on both the Workspace ONE Access cluster and load balancer. The vRealize Automation load balancer uses SSL passthrough, so the certificate is applied only on the vRealize Automation cluster.

See [Managing certificates and DNS configuration under single-node multi-organization deployments](#) and [Managing certificate and DNS configuration under clustered vRealize Automation deployments](#) for more details.

- 2 Create or import the required multi-domain (SAN) certificates for both Workspace ONE Access and vRealize Automation.

You can create certificates in Lifecycle Manager using the Locker service that enables you to create certificates licenses, and passwords. Alternatively, you can use a CA server or some other mechanism to generate certificates.

If you need to add or create additional tenants, you must recreate and apply your vRealize Automation and Workspace ONE Access tenants.

After you create your certificates, you can apply them within Lifecycle Manager using the Lifecycle Operations feature. You must select the environment and product and then the Replace Certificate option on the righthand menu. Then you can select the product. When you replace a certificate, you must re-trust all associated products in your environment.

You must wait for the certificate to be applied and all services to restart before proceeding to the next step.

See [Managing certificates and DNS configuration under single-node multi-organization deployments](#) and [Managing certificate and DNS configuration under clustered vRealize Automation deployments](#) for more details.

- 3 Apply the Workspace ONE Access SAN certificate on the Workspace ONE Access instance or cluster.

- 4 In vRealize Suite Lifecycle Manager, run the Enable Tenancy wizard to enable multi-tenancy and create an alias for the default master tenant.

Enabling tenancy requires that you create an alias for the provider organization master tenant or default tenant. After you enable tenancy, you can access Workspace ONE Access via the master tenant FQDN.

For example, if the existing Workspace ONE Access FQDN is `idm.example.local` and you create an alias of `default-tenant`, after tenancy is enabled, the Workspace ONE Access FQDN changes to `default-tenant.example.local`, and all clients communicating with Workspace ONE Access would now communicate through `default-tenant.example.local`.

- 5 Apply the vRealize Automation SAN certificates on the vRealize Automation instance or cluster.

You can apply SAN certificates through the Lifecycle Manager Lifecycle Operations service. You need to view the details of the environment and then select Replace Certificates on the right menu. You must wait for the certificate replacement task to complete before adding tenants. As part of certificate replacement, vRealize Automation services will restart.

- 6 In Lifecycle Manager, run the Add Tenants wizard to configure the desired tenants.

You add tenants using the Lifecycle Manager Tenant Management page located under Identity and Tenant Management. You can only add tenants for which you have previously configured certificates and DNS settings.

When creating a tenant, you must designate a tenant administrator and you can select the Active Directory connections for this tenant. Available connections are based on those configured in your default or master tenant. You must also select the product or product instance to which the tenant will be associated.

What to do next

After you create tenants, you can use the Lifecycle Manager Tenant Management page located under Identity and Tenant Management to change or add tenant administrators, add Active Directory directories to the tenant and change product associations for the tenant.

You can also log in to your Workspace ONE Access instance to view and validate your tenant configuration.

Managing certificates and DNS configuration under single-node multi-organization deployments

Multi-organization tenancy vRealize Automation configurations rely on a coordinated configuration between several products, and you must ensure that DNS settings and certificates are configured correctly in order for your multi-organization tenancy configuration to function.

This multi-organization configuration assumes single node deployments for the following components:

- Lifecycle Manager
- Workspace ONE Access Identity Manager
- vRealize Automation

Also, it assumes that you are starting with a default tenant, which is your provider organization, and creating two sub-tenants, called tenant-1 and tenant-2.

You can create and apply certificates using the Locker service in vRealize Suite Lifecycle Manager or you can use another mechanism. Lifecycle manager also enables you to replace or re-trust certificates on vRealize Automation or Workspace ONE Access.

DNS Requirements

You must create both main A type records and CNAME type records for system components as described below.

- Create both main A type records for each system component and for each of the tenants that you will create when you enable multi-tenancy.
- Create multi-tenancy A type records for each of the tenants you will create as well as for the master tenant.
- Create multi-tenancy CNAME type records for each of the tenants you will create, not including the master tenant.

Certificate requirements for single node multi-tenancy deployment

You must create two Subject Alternative Name (SAN) certificates, one for Workspace ONE Access and one for vRealize Automation.

- The vRealize Automation certificate lists the hostname of the vRealize Automation server and the names of the tenants you will create.
- The Workspace ONE Access certificate lists the hostname of the Workspace ONE Access server and the tenant names you are creating.

- If you use dedicated SAN names, certificates must be updated manually when you add or delete hosts or change a hostname. You must also update DNS entries for tenants. As an option to simplify configuration, you can use wildcards for the Workspace ONE Access and vRealize Automation certificates. For example, `*.example.com` and `*.vra.example.com`.

Note vRealize Automation 8.x supports wildcard certificates only for DNS names that match the specifications in the Public Suffix list at <https://publicsuffix.org>. For example, `*.myorg.com` is a valid name while `*.myorg.local` is invalid.

Note that Lifecycle Manager does not create separate certificates for each tenant. Instead it creates a single certificate with each tenant hostname listed. For basic configurations, the tenant's CNAME uses the following format: `tenantname.vrahostname.domain`. For high availability configurations, the name uses the following format: `tenantname.vraLBhostname.domain`.

Summary

The following table summarizes DNS and certificate requirements for a single node Workspace ONE Access and single node vRealize Automation deployment.

| DNS Requirements | SAN Certificate Requirements |
|---|---|
| Main A Type Records <code>lcm.example.local</code> <code>WorkspaceOne.example.local</code> <code>vra.example.local</code> | Workspace One Certificate Host Name: <code>WorkspaceOne.example.local</code> , <code>default-tenant.example.local</code> , <code>tenant-1.vra.example.local</code> , <code>tenant-2.vra.example.local</code> |
| Multi-tenancy A Type Records <code>default-tenant.example.local</code> <code>tenant-1.example.local</code> <code>tenant-2.example.local</code> | |
| Multi-Tenancy CNAME Type Records <code>tenant-1.vra.example.local</code> <code>tenant-2.vra.example.local</code> | vRealize Automation Certificate Host Name: <code>vra.example.local</code> , <code>tenant-1.vra.example.local</code> , <code>tenant-2.vra.example.local</code> |

Managing certificate and DNS configuration under clustered vRealize Automation deployments

You must coordinate the certificate and DNS configuration between all applicable components to set up a multi-organization clustered vRealize Automation deployment.

In a typical clustered configuration, there are three Workspace ONE Access appliances and three vRealize Automation appliances as well as a single Lifecycle Manager appliance.

This configuration assumes clustered deployments for the following components:

- Workspace ONE Access Identity Manager appliances:
 - `idm1.example.local`
 - `idm2.example.local`

- idm3.example.local
- idm-lb.example.local
- vRealize Automation appliances:
 - vra-1.example.local
 - vra-2.example.local
 - vra-3.example.local
 - vra-lb.example.local
- Lifecycle Manager appliance

DNS Requirements

You must create both main A type records for each component and for each of the tenants that you will create when you enable multi-tenancy. In addition, you must create multi-tenancy CNAME type records for each of the tenants you will create, not including the master tenant. Finally, you must also create Main A Type records for the Workspace ONE Access and vRealize Automation load balancers.

- Create A type records for the three Workspace ONE Access appliances, and for the vRealize Automation appliances that point to their respective FQDNs.
- In addition, create A type records for the Workspace ONE Access load balancer and the vRealize Automation load balancer that point to their respective FQDNs.
- Create multi-tenancy A Type records for the default tenant and for tenant-1 and tenant-2 that point to the IP address of the Workspace ONE Access load balancer.
- Create CNAME records for tenant-1 and tenant-2 that point to the IP address of the vRealize Automation load balancer.

Subject Alternative Name (SAN) Certificate Requirements

You must create two Workspace ONE Access certificates, one that applies on the cluster appliances and one that applies on the load balancer. In addition, create a certificate that applies to the vRealize Automation appliances, the tenants you are creating, excluding the default tenant, and the load balancer.

- Create a certificate for the Workspace ONE Access appliances that list the FQDNs of the Workspace ONE Access appliances as well as the default tenant and other tenants you create. This certificate should include the IP addresses of the Workspace ONE Access appliances.
- As a best practice, create an SSL termination on the load balancer. To support this termination, create a certificate for the Workspace ONE Access load balancer that lists the FQDN of the Workspace ONE Access load balancer as well as the default tenant and all other tenants you create. This certificate should include the IP address of the load balancer.

- You must create a certificate for vRealize Automation that lists the host names of the three vRealize Automation appliances as well as the related load balancer and the tenants you are creating. In addition, it should list the IP addresses of the three vRealize Automation appliances.
- As an option, to simplify configuration, you can use wildcards for the Workspace ONE Access and vRealize Automation certificates. For example, `*.example.com`, `*.vra.example.com`, and `*.vra-lb.example.com`.

Note vRealize Automation 8.x supports wildcard certificates only for DNS names that match the specifications in the Public Suffix list at <https://publicsuffix.org>. For example, `*.myorg.com` is a valid name while `*.myorg.local` is invalid.

If you are using a clustered Workspace ONE Access configuration, note that Lifecycle Manager cannot update the load balancer certificates, so you must update them manually. Also, if you need to re-register products or services that are external to Lifecycle Manager, this is a manual process.

Summary of DNS entries and certificates for a clustered multi-organization configuration

The following tables outlines DNS Main A Type Records and C Name Type records and certificate requirements for a clustered Workspace ONE Access and clustered vRealize Automation multi-organization deployment.

| DNS Requirements | SAN Certificate Requirements |
|--|---|
| <p>Main A Type Records</p> <ul style="list-style-type: none"> ■ lcm.example.local ■ WorkspaceOne-1.example.local ■ WorkspaceOne-2.example.local ■ WorkspaceOne-3.example.local ■ WorkspaceOne-lb.example.local ■ vra-1.example.local ■ vra-2.example.local ■ vra-3.example.local ■ vra-lb.example.local | <p>Workspace One Certificate</p> <p>Host Name:</p> <ul style="list-style-type: none"> ■ WorkspaceOne-1.example.local ■ WorkspaceOne-2.example.local ■ WorkspaceOne-3.example.local ■ default-tenant.example.local ■ tenant-1.example.local ■ tenant-2.example.local |
| <p>Multi-Tenancy A Type Records</p> <ul style="list-style-type: none"> ■ default-tenant.example.local ■ tenant-1.vra.example.local ■ tenant-2.vra.example.local <p>Note All of the multi-tenancy A Type records must point to the vIDM/WS1A load balancer IP address.</p> | <p>Workspace One LB Certificate (LB Terminated)</p> <p>Host Name:</p> <ul style="list-style-type: none"> ■ WorkspaceOne-lb.example.local ■ default-tenant.example.local ■ tenant-1.example.local ■ tenant-2.example.local |
| <p>Multi-Tenancy CNAME Type Records</p> <ul style="list-style-type: none"> ■ tenant-1.vra-lb.example.local - vra-lb.example.local ■ tenant-2.vra-lb.example.local - vra-lb.example.local | <p>vRealize Automation Certificate</p> <p>Host Name:</p> <ul style="list-style-type: none"> ■ vra-1.example.local ■ vra-2.example.local ■ vra-3.example.local ■ vra-lb.example.local ■ tenant-1.example.local ■ tenant-2.example.local <p>No certificate is required on the vRealize Automation load balancer as it uses SSL passthrough.</p> |

Note Each additional tenant that you add must be listed separately in the vRealize Automation Certificate, Multi-tenancy CNAME records, Multi-tenancy Type A records, Workspace One Certificate and Workspace One LB Certificate.

Note The *.local file names are for example use only. They may not be applicable to most business environments.

Logging in to tenants and adding users in vRealize Automation

After you have created tenants for vRealize Automation in Lifecycle Manager, you can log in to Workspace ONE Access to view your tenants and add users.

You can view tenants created for a vRealize Automation deployment by logging in to the associated Workspace ONE Access instance. The URL to use is `https://default-tenantname.domainname.local` or, for a non-clustered deployment, `https://idm.domainname.local` which will direct you back to the default tenant Workspace ONE Access URL.

You can validate specific tenants in Workspace ONE Access using the following URL: `https://tenant-1.domainname.local`. This URL opens a page that show the users for the specified tenant. You can click **Add User** to create additional users on an ad-hoc basis.

Authorized users can log in to the main provider organization in vRealize Automation using `https://vra.domainname.local`. This view provides access to all vRealize Automation related services.

Authorized users can log in to applicable tenants in vRealize Automation using `https://tenantname.vra.domainname.local`.

For more information about managing users in Workspace ONE Access, see [Managing users and groups](#).

Adding local users

You can add local users to your deployment using the associated Workspace ONE Access instance. Local users are users that are not stored in any external identity provider.

Using vRealize Orchestrator with vRealize Automation multi-organization deployments

You can use vRealize Orchestrator with vRealize Automation multi-organization tenancy deployments.

The default tenant supports integration with the embedded vRealize Orchestrator integration out of the box. vRealize Orchestrator is available pre-configured on the Integrations page of the default tenant. Subtenants do not have any pre-registered vRealize Orchestrator integration. They do have several options to add a vRealize Orchestrator integration.

- Subtenants can add an integration with the embedded vRealize Orchestrator by navigating to **Infrastructure > Connections > Integrations**.

Note If the embedded vRealize Orchestrator is added as an integration to multiple tenants, all the vRealize Orchestrator content, including the plug-in inventory, is shared among these tenants.

- Subtenants can add an external vRealize Orchestrator instance that uses the multi-organization vRealize Automation as an Auth Provider.

Any vRealize Orchestrator instance that uses a vRealize Automation multi-organization deployment as an Auth Provider can be registered to any of the tenants by creating a new integration and providing the vRealize Orchestrator FQDN without providing any credentials.

Working with logs in vRealize Automation

5

You can use the supplied `vracli` command line utility to create and use logs in vRealize Automation.

You can use logs directly in vRealize Automation or you can instead forward all logs to vRealize Log Insight.

This chapter includes the following topics:

- [How do I work with logs and log bundles in vRealize Automation](#)
- [How do I configure log forwarding to vRealize Log Insight in vRealize Automation](#)
- [How do I create or update a syslog integration in vRealize Automation](#)
- [How do I work with content packs](#)

How do I work with logs and log bundles in vRealize Automation

Various services generate logs automatically. You can generate log bundles in vRealize Automation. You can also configure your environment to send logs to vRealize Log Insight.

Use the `--help` argument in the `vracli` command line (for example, `vracli log-bundle --help`) for information about the `vracli` command line utility.

For related information about using vRealize Log Insight, see [How do I configure log forwarding to vRealize Log Insight in vRealize Automation](#).

Log bundle commands

You can create a log bundle to contain all the logs that are generated by the services that you run. A log bundle contains all your service logs. You can use a log bundle for troubleshooting.

In a clustered environment (high availability mode), run the `vracli log-bundle` command on only one node. Logs are pulled from all nodes in the environment. However, in the event of a networking or other cluster issue, logs are pulled from as many nodes as can be reached. For example, if one node is disconnected in a cluster of three nodes, logs are only collected from the two healthy nodes. Output from the `vracli log-bundle` command contains information about any issues found and their workaround steps.

- To create a log bundle, SSH to any node and run the following `vracli` command:

```
vracli log-bundle
```

- To change the timeout value for collecting logs from each node, run the following `vracli` command:

```
vracli log-bundle --collector-timeout $CUSTOM_TIMEOUT_IN_SECONDS
```

For example, if your environment contains large log files, slow networking, or high CPU usage, you can set the timeout to greater than the 1000 second default value.

- To determine the disk space being consumed by a specific service log such as `ebs` or `vro`, run the following `vracli` command and examine the command output:

```
vracli disk-mgr
```

- To configure other options, such as assembly timeout and buffer location, use the following `vracli help` command:

```
vracli log-bundle --help
```

Log bundle structure

The log bundle is a timestamped tar file. The name of the bundle matches the pattern `log-bundle-<date>T<time>.tar` file, for example `log-bundle-20200629T131312.tar`. Typically the log bundle contains logs from all nodes in the environment. In case of an error, it contains as many logs as possible. It minimally contains logs from the local node.

The log bundle consists of the following content:

- Environment file

The environment file contains the output of various Kubernetes maintenance commands. It supplies information about current resource usage per nodes and per pods. It also contains cluster information and description of all available Kubernetes entities.

- Host logs and configuration

The configuration of each host (for example, its `/etc` directory) and the host-specific logs (for example, `journald`) are collected in one directory for each cluster node or host. The directory name matches the host name of the node. The internal contents of the directory match the file system of the host. The number of directories matches the number of cluster nodes.

- Services logs

Logs for Kubernetes services are located in the following folder structure:

- `<hostname>/services-logs/<namespace>/<app-name>/file-logs/<container-name>.log`
- `<hostname>/services-logs/<namespace>/<app-name>/console-logs/<container-name>.log`

An example filename is `my-host-01/services-logs/prelude/vco-app/file-logs/vco-server-app.log`.

- *hostname* is the host name of the node on which the application container is or was running. Typically, there is one instance for each node for each service. For example, 3 nodes = 3 instances.
 - *namespace* is the Kubernetes namespace in which the application is deployed. For user-facing services, this value is `prelude`.
 - *app-name* is the name of the Kubernetes application that produced the logs (for example, `provisioning-service-app`).
 - *container-name* is the name of the container that produced the logs. Some apps consist of multiple containers. For example, the `vco-app` container includes the `vco-server-app` and `vco-controlcenter-app` containers.
- (Legacy) Pod logs

Prior to the logging architecture changes made in vRealize Automation 8.2, services logs were located in each pod's directory in the log bundle. While you can continue to generate pod logs in the bundle by using the `vracli log-bundle --include-legacy-pod-logs` command, doing so is not advised as all log information already resides in each services' logs. Including pod logs can unnecessarily increase the time and space required to generate the log bundle.

Reducing the size of the log bundle

To generate a smaller log bundle, use either of the following `vracli log-bundle` commands:

- `vracli log-bundle --since-days n`

Use this command to collect only the log files that were generated over the past number of days. Otherwise, logs are retained and collected for the past 2 days. For example:

```
vracli log-bundle --since-days 1
```

- `vracli log-bundle --services service_A,service_B,service_C`

Use this command to collect only the logs for the named provided services. For example:

```
vracli log-bundle --services ebs-app,vco-app
```

- `vracli log-bundle --skip-heap-dumps`

Use this command to exclude all heap dumps from the generated log bundle.

Displaying logs

You can output the logs of a service pod or app by using the `vracli logs <pod_name>` command.

The following command options are available:

- `--service`

Displays a merged log for all nodes of the app instead of a single pod

Example: `vracli logs --service abx-service-app`

- `--tail n`

Displays the last *n* lines of the log. The default *n* value is 10.

Example: `vracli logs --tail 20 abx-service-app-8598fcd4b4-tjwhk`

- `--file`

Displays only the specified file. If a file name is not provided, all files are shown.

Example: `vracli logs --file abx-service-app.log abx-service-app-8598fcd4b4-tjwhk`

Understanding log rotation

Regarding log rotation, recognize the following service log considerations:

- All services produce logs. Service logs are stored in a dedicated `/var/log/services-logs` disk.
- All logs are rotated regularly. Rotation occurs either hourly or when a certain size limit is reached.
- All old log rotations are eventually compressed.
- There is no per-service quota for log rotations.
- The system retains as many logs as possible. Automation regularly checks the used disk space for logs. When the space becomes 70% full, older logs are purged until the disk space for logs reaches 60% full.
- You can resize your logs disk if you need more space. See [Increase vRealize Automation appliance disk space](#).

To check the logs disk space, run the following `vracli` commands. The free space of `/dev/sdc (/var/log)` should be near 30% or more for each node.

```
# vracli cluster exec -- bash -c 'current_node; vracli disk-mgr; exit 0'
sc1-10-182-1-103.eng.vmware.com
/dev/sda4(/):
  Total size: 47.80GiB
  Free: 34.46GiB(72.1%)
  Available(for non-superusers): 32.00GiB(66.9%)
  SCSI ID: (0:0)
/dev/sdb(/data):
  Total size: 140.68GiB
  Free: 116.68GiB(82.9%)
```

```

    Available (for non-superusers) : 109.47GiB (77.8%)
    SCSI ID: (0:1)
/dev/sdc (/var/log) :
    Total size: 21.48GiB
    Free: 20.76GiB (96.6%)
    Available (for non-superusers) : 19.64GiB (91.4%)
    SCSI ID: (0:2)
/dev/sdd (/home) :
    Total size: 29.36GiB
    Free: 29.01GiB (98.8%)
    Available (for non-superusers) : 27.49GiB (93.7%)
    SCSI ID: (0:3)

```

How do I configure log forwarding to vRealize Log Insight in vRealize Automation

You can forward logs from vRealize Automation to vRealize Log Insight to take advantage of more robust log analysis and report generation.

vRealize Automation is bundled with a [fluentd-based](#) logging agent. The agent collects and stores logs so that they can be included in a log bundle and examined later. You can configure the agent to forward a copy of the logs to a vRealize Log Insight server by using the vRealize Log Insight REST API. The API allows other programs to communicate with vRealize Log Insight.

For more information about vRealize Log Insight, including documentation for the vRealize Log Insight REST API, see [vRealize Log Insight documentation](#).

Configure the logging agent to continuously forward vRealize Automation logs to vRealize Log Insight by using the supplied `vracli` command line utility.

All log lines are tagged with a host name and environment tag and can be examined in vRealize Log Insight. In a high availability (HA) environment, logs are tagged with different host names, depending on the node that they originated on. The environment tag is configurable by using the `--environment ENV` option as described below in the *Configure or update integration of vRealize Log Insight* section. In an HA environment, the environment tag has the same value for all log lines, regardless of the node they originated on.

Information about how to use the `vracli` command line utility is available by using the `--help` argument in the `vracli` command line. For example: `vracli vrli --help`. For a user-friendly response, begin the command with `vracli -j vrli`.

Note You can only configure a single remote logging integration. vRealize Log Insight is prioritized in the event that both a vRealize Log Insight server and a syslog server are available.

Check existing configuration of vRealize Log Insight

Command

```
vracli vrli
```

Arguments

There are no command line arguments.

Output

The current configuration for vRealize Log Insight integration is output in JSON format.

Exit codes

The following exit codes are possible:

- 0 - Integration with vRealize Log Insight is configured.
- 1 - An exception occurred as part of command execution. Examine the error message for details.
- 61 (ENODATA) - Integration with vRealize Log Insight is not configured. Examine the error message for details.

Example - check integration configuration

```
$ vracli vrli
No vRLI integration configured

$ vracli vrli
{
  "agentId": "0",
  "environment": "prod",
  "host": "my-vrli.local",
  "port": 9543,
  "scheme": "https",
  "sslVerify": false
}
```

Configure or update integration of vRealize Log Insight

Command

```
vracli vrli set [options] FQDN_OR_URL
```

Note After you run the command, it can take up to 2 minutes for the logging agent to apply your specified configuration.

Arguments

- FQDN_OR_URL

Specifies the FQDN or URL address of the vRealize Log Insight server to use for posting logs. Port 9543 and https are used by default. If any of these settings must be changed, you can use a URL instead.

```
vracli vrli set <options> https://FQDN:9543
```

Note You can set a different host scheme (default is https) and port (default for https is 9543, default for http is 9000) to use for sending the logs, as shown in the following samples:

```
vracli vrli set https://HOSTNAME:9543
vracli vrli set --insecure HOSTNAME
vracli vrli set http://HOSTNAME:9000
```

Ports 9543 for https and 9000 for http are used by the vRealize Log Insight ingestion REST API as described in the *Administering vRealize Log Insight* topic *Ports and External Interfaces* in [vRealize Log Insight documentation](#).

■ Options

■ --agent-id SOME_ID

Sets the id of the logging agent for this appliance. The default is 0. Used to identify the agent when posting logs to vRealize Log Insight by using the vRealize Log Insight REST API.

■ --environment ENV

Sets an identifier for the current environment. It will be available in vRealize Log Insight logs as a tag for each log entry. The default is `prod`.

■ --ca-file /path/to/server-ca.crt

Specifies a file that contains the certificate of the certificate authority (CA) that was used to sign the certificate of the vRealize Log Insight server. This forces the logging agent to trust the specified CA and enable it to verify the certificate of the vRealize Log Insight server if it was signed by an untrusted authority. The file may contain a whole certificate chain to verify the certificate. In the case of a self-signed certificate, pass the certificate itself.

■ --ca-cert CA_CERT

Definition is identical to that of `--ca-file` as above, but instead passes the certificate (chain) inline as string.

■ --insecure

Deactivates SSL verification of the server certificate. This forces the logging agent to accept any SSL certificate when posting logs.

■ Advanced options

■ --request-max-size BYTES

Multiple log events are ingested with a single API call. This argument controls the maximum payload size, in bytes, for each request. Valid values are between 4000 and 4000000. The default value is 256000. For related information for allowed values, see vRealize Log Insight events ingestion in the vRealize Log Insight REST API documentation. Setting this value too low can cause logging events that are larger than the allowed size to be dropped.

- `--request-timeout SECONDS`

A call to the API can hang for a number of reasons including problems with the remote, networking issues, and so on. This parameter controls the number of seconds wait for each operation to complete, such as opening a connection, writing data, or awaiting a response, before the call is recognized as failed. The value cannot be less than 1 second. The default is 30.

- `--request-immediate-retries RETRIES`

Logs are buffered in aggregated chunks before they are sent to vRealize Log Insight (see `--buffer-flush-thread-count` below). If an API request fails, the log is retried immediately. The default number of immediate retries is 3. If none of the retries is successful, then the whole log chunk is rolled back and is retried again later.

- `--request-http-compress`

To lower network traffic volumes, you can apply gzip compression to requests that are sent to the vRealize Log Insight server. If this parameter is not specified, no compression is used.

- `--buffer-flush-thread-count THREADS`

For better performance and to limit networking traffic, logs are buffered locally in chunks before they are flushed and sent to the log server. Each chunk contains logs from a single service. Depending on your environment, chunks can grow large and time-consuming to flush. This argument controls the number of chunks that can be flushed simultaneously. The default is 2.

Note When configuring integration over https, if the vRealize Log Insight server is configured to use an untrusted certificate such as a self-signed certificate or a certificate that was signed by an untrusted authority, you must use one of the `--ca-file`, `--ca-cert` or `--insecure` options or the logging agent fails to validate the server identity and does not send logs. When using `--ca-file` or `--ca-cert`, the vRealize Log Insight server certificate must be valid for the server's host name. In all cases, verify the integration by allowing a few minutes for processing and then checking that vRealize Log Insight received the logs.

Output

No output is expected.

Exit codes

The following exit codes are possible:

- 0 - The configuration was updated.
- 1 - An exception occurred as part of the execution. Examine the error message for details.

Examples - Configure or update integration configuration

The following example statements are shown in separate command lines, however the arguments can be combined in a single command line. For example, you can include multiple arguments when using `vracli vrli set {somehost}` Or `vracli vrli set --ca-file path/to/server-ca.crt` to modify the default agent ID or environment values. For related information, see the online command help at `vracli vrli --help`.

```
$ vracli vrli set my-vrli.local
$ vracli vrli set 10.20.30.40
$ vracli vrli set --ca-file /etc/ssl/certs/ca.crt 10.20.30.40
$ vracli vrli set --ca-cert "$(cat /etc/ssl/certs/ca.crt)" 10.20.30.40
$ vracli vrli set --insecure http://my-vrli.local:8080
$ vracli vrli set --agent-id my-vrli-agent my-vrli.local
$ vracli vrli set --request-http-compress
$ vracli vrli set --environment staging my-vrli.local
$ vracli vrli set --environment staging --request-max-size 10000 --request-timeout 120 --
request-immediate-retries 5 --buffer-flush-thread-count 4 my-vrli.local
```

Clear integration of vRealize Log Insight

Command

```
vracli vrli unset
```

Note After you run the command, it can take up to 2 minutes for the logging agent to apply your specified configuration.

Arguments

There are no command line arguments.

Output

Confirmation is output in plain text format.

Exit codes

The following exit codes are available:

- 0 - The configuration was cleared or no configuration existed.
- 1 - An exception occurred as part of the execution. Examine the error message for details.

Examples - Clear integration

```
$ vracli vrli unset
Clearing vRLI integration configuration

$ vracli vrli unset
No vRLI integration configured
```

How do I create or update a syslog integration in vRealize Automation

You can configure vRealize Automation to send your logging information to remote syslog servers.

The `vracli remote-syslog set` command is used to create a syslog integration or overwrite existing integrations.

vRealize Automation remote syslog integration supports the following connection types:

- Over UDP.
- Over TCP without TLS.

Note To create a syslog integration without using TLS, add the `--disable-ssl` flag to the `vracli remote-syslog set` command.

- Over TCP with TLS.

Note You can only configure a single remote logging integration. vRealize Log Insight is prioritized in the event that both a vRealize Log Insight server and a syslog server are available.

For information on configuring logging integration with vRealize Log Insight, see [How do I configure log forwarding to vRealize Log Insight in vRealize Automation](#).

Prerequisites

Configure a remote syslog server.

Procedure

- 1 Log in to the vRealize Automation appliance command line as **root**.

- 2 To create an integration to a syslog server, run the `vracli remote-syslog set` command.

```
vracli remote-syslog set -id name_of_integration protocol_type://
syslog_URL_or_FQDN:syslog_port
```

Note If you do not enter a port in the `vracli remote-syslog set` command, the port value defaults to 514.

Note You can add a certificate to the syslog configuration. To add a certificate file, use the `--ca-file` flag. To add a certificate as plaintext, use `--ca-cert` flag.

- 3 (Optional) To overwrite an existing syslog integration, run the `vracli remote-syslog set` and set the `-id` flag value to the name of the integration you want to overwrite.

Note By default, the vRealize Automation appliance requests that you confirm that you want to overwrite the syslog integration. To skip the confirmation request, add the `-f` or `--force` flag to the `vracli remote-syslog set` command.

What to do next

To review the current syslog integrations in the appliance, run the `vracli remote-syslog` command.

How do I delete a syslog integration for logging in vRealize Automation

You can delete syslog integrations from your vRealize Automation appliance by running the `vracli remote-syslog unset` command.

Prerequisites

Create one or more syslog integrations in the vRealize Automation appliance. See [How do I create or update a syslog integration in vRealize Automation](#).

Procedure

- 1 Log in to the vRealize Automation appliance command line as **root**.
- 2 Delete syslog integrations from the vRealize Automation appliance using either of the following methods:
 - To delete a specific syslog integration, run the `vracli remote-syslog unset -id Integration_name` command.

- To delete all syslog integrations on the vRealize Automation appliance, run the `vracli remote-syslog unset` command without the `-id` flag.

Note By default, the vRealize Automation appliance requests that you confirm that you want to delete all syslog integrations. To skip the confirmation request, add the `-f` or `--force` flag to the `vracli remote-syslog unset` command.

How do I work with content packs

Content packs are hosted in Log Insight and contain dashboards, extracted fields, saved queries, and alerts that are related to a specific product or set of logs. You can install community supported content packs from the VMware Sample Exchange and other content packs from the Content Pack Marketplace.

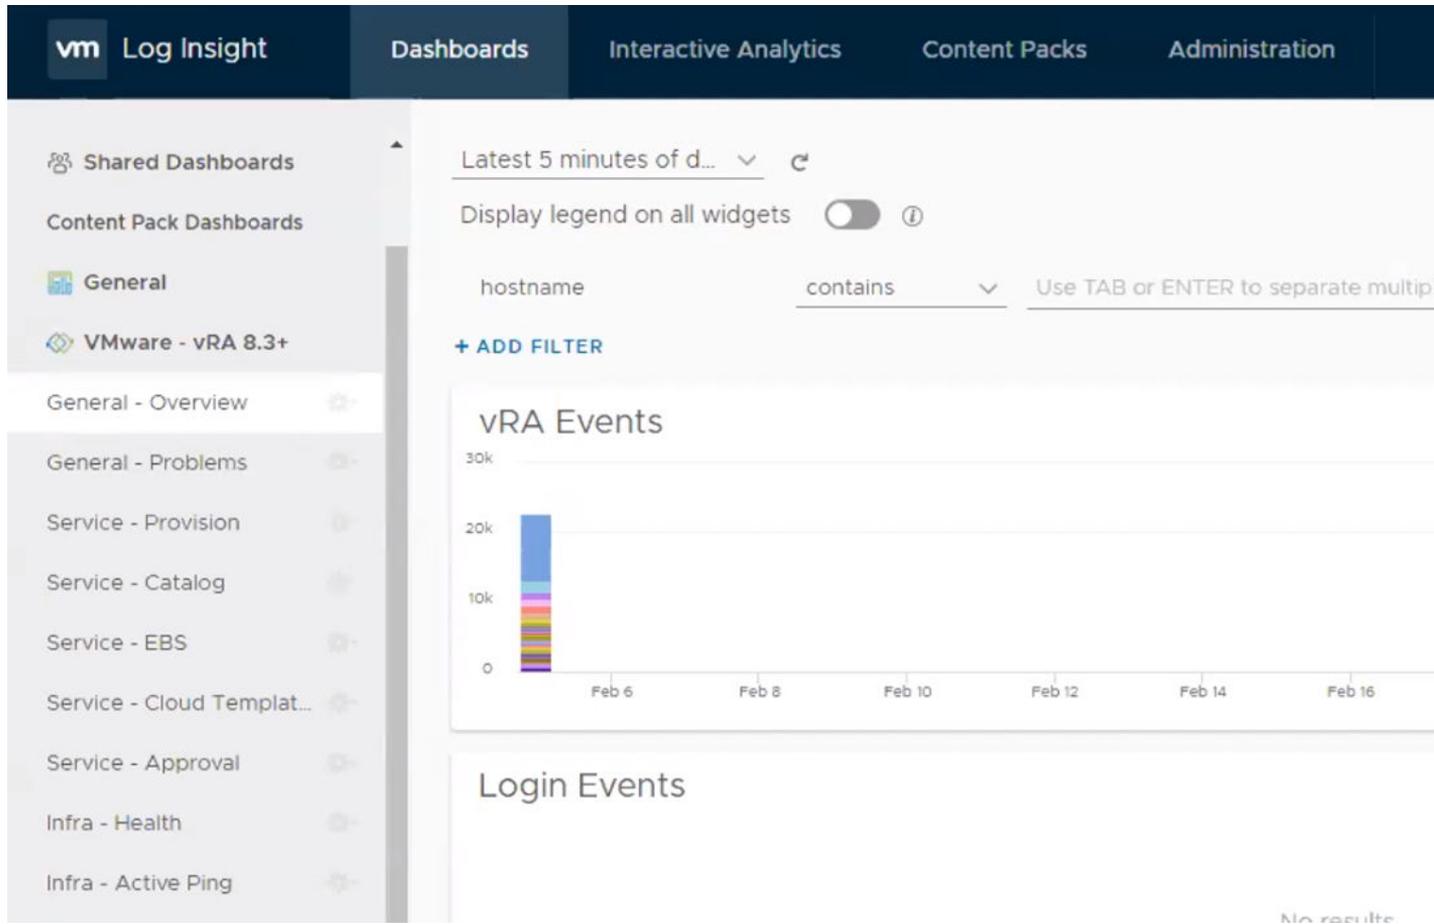
VMware vRealize Log Insight delivers automated log management through aggregation, analytics and search, enabling operational intelligence and enterprise-wide visibility in dynamic hybrid cloud environments. Content packs are plugins to VMware vRealize Log Insight that provide pre-defined knowledge about specific types of events such as log messages.

To download a content pack, from Log Insight navigate to **Content Packs > Marketplace**. You can also import content packs by clicking **+ Import Content Pack**.

vRA 8.x Content Pack

The VMware vRealize Automation content pack provides a consolidated summary of log events across all vRA environment components. It includes several dashboards that provide a general overview, insight on errors and operations, and overall health of your vRA instance. These dashboards are listed on the **Dashboard** tab along with all other Log Insight dashboards. Once loaded, it can take up to 30 seconds for the dashboards to populate with metrics.

Note You cannot upgrade from vRA 7.5+ content pack to the vRA 8.3 content pack. You must install the vRA 8.3 content pack. Once installed, the 8.3 and 7.5 content packs work separately.

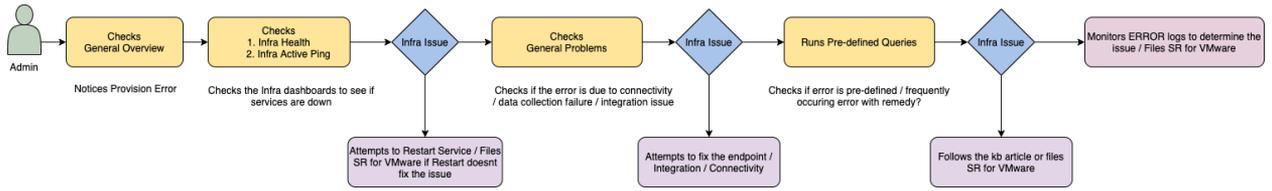


The vRealize Automation content pack includes these dashboards:

- General - Overview: Captures an overview of high level metrics for vRA.
- General - Problems:
- Service - Provision: Captures issues related to the provisioning service.
- Service - Catalog: Captures issues related to the catalog service.
- Service - EBS: Captures issues related to the event broker service.
- Service - Cloud Templates: Captures errors and metrics related to Cloud Assembly cloud templates, custom resources, and resource actions.
- Service - Approval: Captures errors and metrics related to approvals.
- Infra - Health: Captures when pods are restarted over time. This dashboard is essential to detect outages dues to resource limits.
- Infra - Active Ping: Captures the health check url over time.

Each dashboard contains individual widgets that provide more focused analytics. To view what type of analysis is performed in each widget, click the information  icon.

As a vRealize Automation administrator, you can follow this general content pack workflow to identify errors and troubleshoot.



For more information on the vRealize Automation 8.3 Content Pack, see [vRealize Automation 8.3+ Log Insight Content Pack](#) and [How do I configure log forwarding to vRealize Log Insight](#).

Participating in the Customer Experience Improvement Program for vRealize Automation

6

This product participates in VMware's Customer Experience Improvement Program (CEIP). The CEIP provides VMware with information that allows VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products.

Details regarding the data collected through CEIP, and the purposes for which it is used by VMware, are described on the [Customer Experience Improvement Program](#) page.

This chapter includes the following topics:

- [How do I join or leave the Customer Experience Improvement Program for vRealize Automation](#)
- [How do I configure the data collection time for the Customer Experience Improvement Program for vRealize Automation](#)

How do I join or leave the Customer Experience Improvement Program for vRealize Automation

Join or leave the Customer Experience Improvement Program (CEIP) from the vRealize Automation appliance command line.

You can join the CEIP program when you install vRealize Automation and with the vRealize Lifecycle Manager (LCM). You can also join or leave the program by using command line options after installation.

To join the Customer Experience Improvement Program by using command line options:

- 1 Log in to the vRealize Automation appliance command line as **root**.
- 2 Run the `vracli ceip on` command.
- 3 Review the Customer Experience Improvement Program information and run the `vracli ceip on --acknowledge-ceip` command.
- 4 To restart the vRealize Automation services, run the `/opt/scripts/deploy.sh` command.

To leave the Customer Experience Improvement Program by using command line options:

- 1 Log in to the vRealize Automation appliance command line as **root**.
- 2 Run the `vracli ceip off` command.

- 3 To restart the vRealize Automation services, run the `/opt/scripts/deploy.sh` command.

How do I configure the data collection time for the Customer Experience Improvement Program for vRealize Automation

You can set the day and time when the Customer Experience Improvement Program (CEIP) sends data to VMware.

Procedure

- 1 Log in to the vRealize Automation appliance command line as **root**.

- 2 Open the following file in a text editor.

```
/etc/telemetry/telemetry-collector-vami.properties
```

- 3 Edit the properties for day of week (dow) and hour of day (hod).

| Property | Description |
|--|--|
| <code>frequency.dow=<day-of-week></code> | Day when data collection occurs. |
| <code>frequency.hod=<hour-of-day></code> | Local time of day when data collection occurs. Possible values are 0–23. |

- 4 Save and close `telemetry-collector-vami.properties`.

- 5 Apply the settings by entering the following command.

```
vcac-config telemetry-config-update --update-info
```

Changes are applied to all nodes in your deployment.

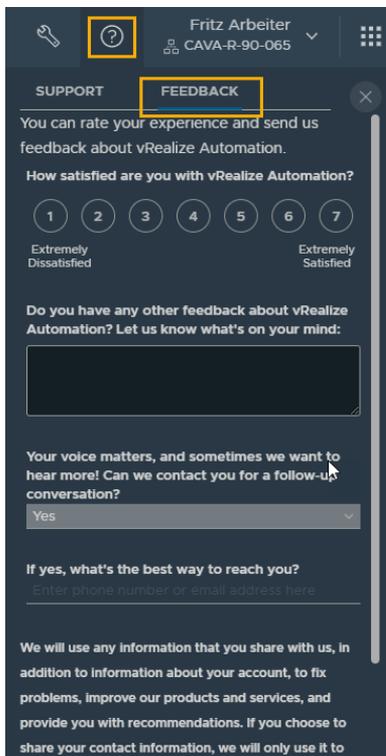
Turning on the in-product feedback form in vRealize Automation

7

You can enable your users to provide feedback to the vRealize Automation development team. Your feedback is important to our development process.

What is the feedback form

The feedback form is located in the support panel on a tab labeled Feedback. To open the form, click the **Help** button and then click **Feedback**.



How do I make the feedback form available to my users

The feedback form requires that your vRealize Automation host has Internet access and that the following base URLs are included in your Internet allowlist.

- <https://lumos.vmware.com/>

- <https://feedback.esp.vmware.com/>

If the host does not have Internet access, the form is not available in the Help pane.