VMware Cloud Foundation 4.2.1 Release Notes

VMware Cloud Foundation 4.2.1

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

Contents

Introduction

1

VMware Cloud Foundation 4.2.1 | 25 MAY 2021 | Build 18016307

Check for additions and updates to these release notes.

What's New

2

The VMware Cloud Foundation (VCF) 4.2.1 release includes the following:

- **Security fixes for Photon OS:** SDDC Manager 4.2.1 contains security updates for Photon OS packages from PHSA-2021-3.0-185 to PHSA-2021-3.0-209. To view information about the updates, see Photon OS release 3.0 advisories.
- **Security fixes for VMware vCenter Server Appliance**: See VMSA-2021-0010.
- **BOM updates**: Updated Bill of Materials with new product versions.

VMware Cloud Foundation Bill of Materials (BOM)

3

The Cloud Foundation software product is comprised of the following software Bill-of-Materials (BOM). The components in the BOM are interoperable and compatible.

VMware Response to Apache Log4j Remote Code Execution Vulnerability: VMware Cloud Foundation is impacted by CVE-2021-44228, and CVE-2021-45046 as described in VMSA-2021-0028. To remediate these issues, see Workaround instructions to address CVE-2021-44228 & CVE-2021-45046 in VMware Cloud Foundation (KB 87095).

Software Component	Version	Date	Build Number
Cloud Builder VM	4.2.1	25 MAY 2021	18016307
SDDC Manager	4.2.1	25 MAY 2021	18016307
VMware vCenter Server Appliance	7.0.1.00301	25 MAY 2021	17956102
VMware ESXi	7.0 Update 1d	04 FEB 2021	17551050*
VMware NSX-T Data Center	3.1.2	17 APR 2021	17883596
VMware vRealize Suite Lifecycle Manager	8.2 Patch 2	04 FEB 2021	17513665
Workspace ONE Access	3.3.4	04 FEB 2021	17498518
vRealize Automation	8.2	06 OCT 2020	16980951
vRealize Log Insight	8.2	06 OCT 2020	16957702
vRealize Log Insight Content Pack for NSX-T	3.9.2	n/a	n/a
vRealize Log Insight Content Pack for Linux	2.1	n/a	n/a
vRealize Log Insight Content Pack for Linux - Systemd	1.0	n/a	n/a
vRealize Log Insight Content Pack for vRealize Suite Lifecycle Manager 8.0.1+	1.0.2	n/a	n/a

Software Component	Version	Date	Build Number
vRealize Log Insight Content Pack for VMware Identity Manager	2.0	n/a	n/a
vRealize Operations Manager	8.2	06 OCT 2020	16949153
vRealize Operations Management Pack for VMware Identity Manager	1.1	n/a	n/a

* VMware ESXi 7.0 Update 1d is a patch release and does not have an ISO available for download on My VMware. You can create an ISO to install the correct version of ESXi on your servers. See Create a Custom ISO Image for ESXi.

- VMware vSAN is included in the VMware ESXi bundle.
- You can use vRealize Suite Lifecycle Manager to deploy vRealize Automation, vRealize Operations Manager, vRealize Log Insight, and Workspace ONE Access using the VMware Validated Design 6.2 documentation.
- vRealize Log Insight content packs are installed when you deploy vRealize Log Insight.
- The vRealize Operations Manager management pack is installed when you deploy vRealize Operations Manager.
- VMware Solution Exchange and the vRealize Log Insight in-product marketplace store only the latest versions of the content packs for vRealize Log Insight. The Bill of Materials table contains the latest versions of the packs that were available at the time VMware Cloud Foundation is released. When you deploy the Cloud Foundation components, it is possible that the version of a content pack within the in-product marketplace for vRealize Log Insight is newer than the one used for this release.

VMware Software Edition License Information

4

The SDDC Manager software is licensed under the Cloud Foundation license. As part of this product, the SDDC Manager software deploys specific VMware software products.

The following VMware software components deployed by SDDC Manager are licensed under the VMware Cloud Foundation license:

- VMware ESXi
- VMware vSAN
- VMware NSX-T Data Center

The following VMware software components deployed by SDDC Manager are licensed separately:

- vCenter Server

  NOTE: Only one vCenter Server license is required for all vCenter Servers deployed in a Cloud Foundation system.

For details about the specific VMware software editions that are licensed under the licenses you have purchased, see the Cloud Foundation Bill of Materials (BOM) section above.

For general information about the product, see VMware Cloud Foundation.

Supported Hardware

5

For details on supported configurations, see the VMware Compatibility Guide (VCG) and the Hardware Requirements section on the Prerequisite Checklist tab in the Planning and Preparation Workbook.

Documentation

6

To access the Cloud Foundation documentation, go to the VMware Cloud Foundation product documentation.

To access the documentation for VMware software products that SDDC Manager can deploy, see the product documentation and use the drop-down menus on the page to choose the appropriate version:

- VMware vSphere product documentation, also has documentation about ESXi and vCenter Server
- VMware vSAN product documentation
- VMware NSX-T Data Center product documentation

Browser Compatibility and Screen Resolutions

7

The Cloud Foundation web-based interface supports the latest two versions of the following web browsers except Internet Explorer:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge
- Internet Explorer: Version 11

For the Web-based user interfaces, the supported standard resolution is 1024 by 768 pixels. For best results, use a screen resolution within these tested resolutions:

- 1024 by 768 pixels (standard)
- 1366 by 768 pixels
- 1280 by 1024 pixels
- 1680 by 1050 pixels

Resolutions below 1024 by 768, such as 640 by 960 or 480 by 800, are not supported.

Installation and Upgrade Information

8

You can install VMware Cloud Foundation 4.2.1 as a new release or upgrade to VMware Cloud Foundation 4.2.1 from VMware Cloud Foundation 4.2, 4.1.0.1, or 4.1.

Installing as a New Release

The new installation process has three phases:

Phase One: Prepare the Environment

The *Planning and Preparation Workbook* provides detailed information about the software, tools, and external services that are required to implement a Software-Defined Data Center (SDDC) with VMware Cloud Foundation, using a standard architecture model.

Phase Two: Image all servers with ESXi

Image all servers with the ESXi version mentioned in the Cloud Foundation Bill of Materials (BOM) section. See the *VMware Cloud Foundation Deployment Guide* for information on installing ESXi.

Phase Three: Install Cloud Foundation 4.2.1

Refer to the *VMware Cloud Foundation Deployment Guide* for information on deploying Cloud Foundation.

Upgrading to VMware Cloud Foundation 4.2.1

You can upgrade to VMware Cloud Foundation 4.2.1 from VMware Cloud Foundation 4.2, 4.1.0.1, or 4.1. For more information see VMware Cloud Foundation Lifecycle Management.

Access to Private APIs

9

Access to private APIs that use basic authentication is being deprecated in an upcoming release. You must switch to using public APIs.

Resolved Issues

10

The following issues are resolved in this release:

- Credential logging vulnerability as described in VMSA-2022-0003. See KB 87050 for more information.
- If you used the Management Plane API to configure L3 forwarding mode, policy will overwrite Management Plane change with default mode IPV4_ONLY after NSX-T is upgraded to 3.1 and will disrupt IPv6 connectivity
- Offline bundle download fails to update LCM 2.0 manifest
- During offline bundle download, LCM manifest file upload to SDDC Manager VM fails with the error "Release BOM cannot be updated after EOL"
- Rotating vCenter Server or ESXi passwords fails

Known Issues

11

This chapter includes the following topics:

VMware Cloud Foundation Known Issues

- **Workload Management does not support NSX-T Data Center Federation**

  You cannot deploy Workload Management (vSphere with Tanzu) to a workload domain when that workload domain's NSX-T Data Center instance is participating in an NSX-T Data Center Federation.

  Workaround: None.

- **NSX-T Guest Introspection (GI) and NSX-T Service Insertion (SI) are not supported on stretched clusters**

  There is no support for stretching clusters where NSX-T Guest Introspection (GI) or NSX-T Service Insertion (SI) are enabled. VMware Cloud Foundation detaches Transport Node Profiles from AZ2 hosts to allow AZ-specific network configurations. NSX-T GI and NSX-T SI require that the same Transport Node Profile be attached to all hosts in the cluster.

  Workaround: None

- **Stretched clusters and Workload Management**

  You cannot stretch a cluster on which Workload Management is deployed.

  Workaround: None

- **Special characters not allowed in the Username, Password, and Template Name fields on the Microsoft CA Configuration page**

  If any of the following special characters are used in the Username, Password, or Template Name fields on the Microsoft CA Configuration page, the configuration cannot be saved:

  - Ampersand (&)
  - Single quote (')
  - Double quote (")
  - Less than (<)
  - Greater than (>)
  - Tab

  Workaround: Delete the special character and then save the configuration.

- **Different vCenter Server build numbers on SDDC Manager and vCenter Server UI**

  The vCenter Server build number on the vCenter Server UI is different from the build number displayed on SDDC Manager.

  Workaround: None.

Upgrade Known Issues

- **Async Patch Tool Known Issues**

  The Async Patch Tool is a utility that allows you to apply critical patches to certain VMware Cloud Foundation components (NSX-T Manager, vCenter Server, and ESXi) outside of VMware Cloud Foundation releases. The Async Patch Tool also allows you to enable upgrade of an async patched system to a new version of VMware Cloud Foundation.

  See the Async Patch Tool Release Notes for known issues.

- **Cluster-level ESXi upgrade fails**

  Cluster-level selection during upgrade does not consider the health status of the clusters and may show a cluster's status as **Available**, even for a faulty cluster. If you select a faulty cluster, the upgrade fails.

  Workaround: Always perform an update precheck to validate the health status of the clusters. Resolve any issues before upgrading.

- **When you skip hosts during an ESXi upgrade of a vLCM-enabled workload domain, the upgrade may fail**

  Due to a known vSphere issue, the ESXi upgrade may fail with a "ConstraintValidationException" error when a host is skipped.

  Workaround: Check the vSphere Client and its logs for details on what caused the error. Resolve the issues and retry the upgrade.

- **When you upgrade to VMware Cloud Foundation 4.1, one of the vSphere Cluster Services (vCLS) agent VMs gets placed on local storage**

  vSphere Cluster Services (vCLS) is new functionality in vSphere 7.0 Update 1 that ensures that cluster services remain available, even when the vCenter Server is unavailable. vCLS deploys three vCLS agent virtual machines to maintain cluster services health. When you upgrade to VMware Cloud Foundation 4.1, one of the vCLS VMs may get placed on local storage instead of shared storage. This could cause issues if you delete the ESXi host on which the VM is stored.

  Workaround: Deactivate and reactivate vCLS on the cluster to deploy all the vCLS agent VMs to shared storage.

  a Check the placement of the vCLS agent VMs for each cluster in your environment.

    1 In the vSphere Client, select **Menu > VMs and Templates**.

    2 Expand the vCLS folder.

    3 Select the first vCLS agent VM and click the Summary tab.

    4 In the Related Objects section, check the datastore listed for Storage. It should be the vSAN datastore. If a vCLS agent VM is on local storage, you need to deactivate vCLS for the cluster and then re-enable it.

    5 Repeat these steps for all vCLS agent VMs.

  b Deactivate vCLS for clusters that have vCLS agent VMs on local storage.

    1 In the vSphere Client, click **Menu > Hosts and Clusters**.

    2 Select a cluster that has a vCLS agent VM on local storage.

    3 In the web browser address bar, note the moref id for the cluster.

      For example, if the URL displays as https://vcenter-1.vrack.vsphere.local/ui/app/cluster;nav=h/urn:vmomi:ClusterComputeResource:domain-c8:503a0d38-442a-446f-b283-d3611bf035fb/summary, then the moref id is domain-c8.

    4 Select the vCenter Server containing the cluster.

    5 Click **Configure > Advanced Settings**.

    6 Click **Edit Settings**.

    7 Change the value for `config.vcls.clusters.<moref id>.enabled` to `false` and click **Save**.

      If the `config.vcls.clusters.<moref id>.enabled` setting does not appear for your moref id, then enter its Name and `false` for the Value and click **Add**.

    8 Wait a couple of minutes for the vCLS agent VMs to be powered off and deleted. You can monitor progress in the Recent Tasks pane.

c Enable vCLS for the cluster to place the vCLS agent VMs on shared storage.

1 Select the vCenter Server containing the cluster and click **Configure > Advanced Settings**.

2 Click **Edit Settings**.

3 Change the value for `config.vcls.clusters.<moref id>.enabled` to `true` and click **Save**.

4 Wait a couple of minutes for the vCLS agent VMs to be deployed and powered on. You can monitor progress in the Recent Tasks pane.

d Check the placement of the vCLS agent VMs to make sure they are all on shared storage

- **You are unable to update NSX-T Data Center in the management domain or in a workload domain with vSAN principal storage because of an error during the NSX-T transport node precheck stage**

  In SDDC Manager, when you run the upgrade precheck before updating NSX-T Data Center, the NSX-T transport node validation results with the following error.

  ```
  No coredump target has been configured. Host core dumps cannot be saved.:System logs on host sfo01-m01-esx04.sfo.rainpole.io are stored on non-persistent storage. Consult product documentation to configure a syslog server or a scratch partition.
  ```

  Because the upgrade precheck results with an error, you cannot proceed with updating the NSX-T Data Center instance in the domain. VMware Validated Design supports vSAN as the principal storage in the management domain. However, vSAN datastores do no support scratch partitions. See VMware KB article 2074026.

  Workaround: Deactivate the update precheck validation for the subsequent NSX-T Data Center update.

  a Log in to SDDC Manager as `vcf` using a Secure Shell (SSH) client.

  b Open the `application-prod.properties` file for editing.

  ```
  vi /opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties
  ```

  c Add the following property and save the file.

  ```
  lcm.nsxt.suppress.prechecks=true
  ```

  d Restart the life cycle management service.

  ```
  systemctl restart lcm
  ```

  e Log in to the SDDC Manager user interface and proceed with the update of NSX-T Data Center.

- **NSX-T upgrade may fail at the step NSX T TRANSPORT NODE POSTCHECK STAGE**

  NSX-T upgrade may not proceed beyond the `NSX T TRANSPORT NODE POSTCHECK STAGE`.

  Workaround: Contact VMware support.

- **Bundle transfer utility command retrieves incorrect bundle types**

  When you run the bundle transfer utility command to retrieve install bundles, the results include install and configuration drift bundles.

  Workaround: Review bundle components to validate the bundle type.

- **Workload Management upgrade fails**

  Workload Management can be upgraded only after NSX-T, vCenter Server, and ESXi have been upgraded. If you try upgrading Workload Management before upgrading these components, the upgrade fails.

  Workaround: Contact VMware Support.

- **ESXi hosts with an HPE custom image cannot be upgraded**

  On ESXi hosts with an HPE custom image, the HPE Agentless Management Service (amsd) blocks the upgrade of NSX components with the following error message:

  ```
  Error loading plugin /usr/lib/vmware/esxcli/ext/smad_rev.xml skipping. Error was:
  Error while trying to register the plugin /usr/lib/vmware/esxcli/ext/smad_rev.xml
  Duplicate top-level namespaces must have the samesma The functionality in /usr/lib/
  vmware/esxcli/ext/smad_rev.xml will not be available until this issue is resolved.
  ```

  Impacted versions and components are listed below.

  HPE Server Line : ProLiant

  HPE Add-on/Custom Image Version : 701.0.0.10.7.0-71

  ESXi Base Image build : 7.0 U1d - 17551050

  amsdComponent Version : 701.11.7.1.3-1

  amshelprComponent Version : 701.11.7.0.14-1

  HPE Server Line : Synergy

  HPE Add-on/Custom Image Version : 701.0.0.10.7.5-16

  ESXi Base Image build : 7.0 U1d - 17551050

  amsdComponent Version : 701.11.7.1.3-1

  amshelprComponent Version : 701.11.7.0.14-1

  Workaround: Before upgrading to VMware Cloud Foundation 4.2.1, apply the fix described in the HPE Customer advisory https://support.hpe.com/hpesc/public/docDisplay?docId=emr_na-a00116792en_us.

Bring-up Known Issues

- **The Cloud Foundation Builder VM remains locked after more than 15 minutes.**

The VMware Imaging Appliance (VIA) locks out the admin user after three unsuccessful login attempts. Normally, the lockout is reset after fifteen minutes but the underlying Cloud Foundation Builder VM does not automatically reset.

Log in to the VM console of the Cloud Foundation Builder VM as the `root` user. Unlock the account by resetting the password of the admin user with the following command:

```
pam_tally2 --user=<user> --reset
```

SDDC Manager Known Issues

- **Deactivating CEIP on SDDC Manager does not deactivate CEIP on vRealize Automation and vRealize Suite Lifecycle Manager**

  When you deactivate CEIP on the SDDC Manager Dashboard, data collection is not deactivated on vRealize Automation and vRealize Suite Lifecycle Manager. This is because of API deprecation in vRealize Suite 8.x.

  Workaround: Manually deactivate CEIP in vRealize Automation and vRealize Suite Lifecycle Manager. For more information, see VMware vRealize Automation Documentation and VMware vRealize Suite Lifecycle Manager Documentation.

- **Generate CSR task for a component hangs**

  When you generate a CSR, the task may fail to complete due to issues with the component's resources. For example, when you generate a CSR for NSX Manager, the task may fail to complete due to issues with an NSX Manager node. You cannot retry the task once the resource is up and running again.

  Workaround:

  a Log in to the UI for the component to troubleshoot and resolve any issues.

  b Using SSH, log in to the SDDC Manager VM with the user name `vcf`.

  c Type **su** to switch to the root account.

  d Run the following command:

  ```
  systemctl restart operationsmanager
  ```

  e Retry generating the CSR.

- **SoS utility options for health check are missing information**

  Due to limitations of the ESXi service account, some information is unavailable in the following health check options:

  - `--hardware-compatibility-report`: No `Devices and Driver` information for ESXi hosts.
  - `--storage-health`: No `vSAN Health Status` or `Total no. of disks` information for ESXi hosts.

  Workaround: None.

- **Host information on the Inventory -> Hosts page takes too long to load or does not load**

  When a large number of hosts are unassigned, the Inventory -> Hosts page may take up to three minutes to reflect hosts information or may display the following message: `Failed to load host details, Please retry or contact the service provider and provide the reference token`

  Workaround: Use VMware Cloud Foundation APIs for all host-related tasks.

- **Host commissioning may fail if API /v1/hosts/validations/commissions is used immediately after the host validation API /v1/hosts/validationsDescription**

  If you run the host validation API `/v1/hosts/validations/commissions` immediately after the host validation API `/v1/hosts/validations`, the validation workflow may delete temporary truststore created by the host commissioning workflow. This may cause host commissioning to fail.

  Workaround: After host validation, wait for at least one minute before running the host commissioning API.

Workload Domain Known Issues

- **Cannot reuse a static IP pool that includes special characters in its name**

  If you chose Static IP Pool as the IP allocation method when creating a VI workload domain and you used special characters or spaces in the IP pool name, you are not able to reuse the IP pool when creating a new VI workload domain or adding a vSphere cluster to the workload domain.

  Workaround: Use only supported characters when naming a static IP pool. Supported characters:

  - a-z
  - A-Z
  - 0-9
  - - and _
  - No spaces

  If you have an existing static IP pool that includes unsupported characters in its name, you can use the NSX Manager UI to rename it.

- **Adding host fails when host is on a different VLAN**

  A host add operation can sometimes fail if the host is on a different VLAN.

  a Before adding the host, add a new portgroup to the VDS for that cluster.

  b Tag the new portgroup with the VLAN ID of the host to be added.

  c Add the Host. This workflow fails at the "Migrate host vmknics to dvs" operation.

d Locate the failed host in vCenter, and migrate the vmk0 of the host to the new portgroup you created in step 1. For more information, see Migrate VMkernel Adapters to a vSphere Distributed Switch in the vSphere product documentation.

e Retry the Add Host operation.

NOTE: If you later remove this host in the future, you must manually remove the portgroup as well if it is not being used by any other host.

- **Deploying partner services on an NSX-T workload domain displays an error**

  Deploying partner services, such as McAfee or Trend, on a workload domain enabled for vSphere Update Manager (VUM), displays the “Configure NSX at cluster level to deploy Service VM” error.

  Workaround: Attach the Transport node profile to the cluster and try deploying the partner service. After the service is deployed, detach the transport node profile from the cluster.

- **If the witness ESXi version does not match with the host ESXi version in the cluster, vSAN cluster partition may occur**

  vSAN stretch cluster workflow does not check the ESXi version of the witness host. If the witness ESXi version does not match the host version in the cluster, then vSAN cluster partition may happen.

  Workaround:

  a Upgrade the witness host manually with the matching ESXi version using the vCenter VUM functionality.

  b Replace or deploy the witness appliance matching with the ESXi version.

- **vSAN partition and critical alerts are generated when the witness MTU is not set to 9000**

  If the MTU of the witness switch in the witness appliance is not set to 9000, the vSAN stretch cluster partition may occur.

  Workaround: Set the MTU of the witness switch in the witness appliance to 9000 MTU.

- **VI workload domain creation or expansion operations fail**

  If there is a mismatch between the letter case (upper or lower) of an ESXi host's FQDN and the FQDN used when the host was commissioned, then workload domain creation and expansion may fail.

  Workaround: ESXi hosts should have lower case FQDNs and should be commissioned using lower case FQDNs.

- **Adding a host to a vLCM-enabled workload domain configured with the Dell Hardware Support Manager (OMIVV) fails**

  When you try to add a host to a vSphere cluster for a workload domain enabled with vSphere Lifecycle Manager (vLCM), the task fails and the domain manager log reports "The host (host-name) is currently not managed by OMIVV." The domain manager logs are located at /var/log/vmware/vcf/domainmanager on the SDDC Manager VM.

Workaround: Update the hosts inventory in OMIVV and retry the add host task in the SDDC Manager UI. See the Dell documentation for information about updating the hosts inventory in OMIVV.

- **Adding a vSphere cluster or adding a host to a workload domain fails**

  Under certain circumstances, adding a host or vSphere cluster to a workload domain fails at the `Configure NSX-T Transport Node or Create Transport Node Collection` subtask.

  Workaround:

  a Enable SSH for the NSX Manager VMs.

  b SSH into the NSX Manager VMs as `admin` and then log in as `root`.

  c Run the following command on each NSX Manager VM:

  ```
  sysctl -w net.ipv4.tcp_en=0
  ```

  d Login to NSX Manager UI for the workload domain.

  e Navigate to **System > Fabric > Nodes > Host Transport Nodes**.

  f Select the vCenter server for the workload domain from the **Managed by** drop-down menu.

  g Expand the vSphere cluster and navigate to the transport nodes that are in a `partial success` state.

  h Select the check box next to a `partial success` node, click **Configure NSX.** .

  i Click `Next` and then click `Apply`.

  j Repeat steps 7-9 for each `partial success` node.

  When all host issues are resolved, transport node creation starts for the failed nodes. When all hosts are successfully created as transport nodes, retry the failed add vSphere cluster or add host task from the SDDC Manager UI.

- **The vSAN Performance Service is not enabled for vSAN clusters when CEIP is not enabled**

  If you do not enable the VMware Customer Experience Improvement Program (CEIP) in SDDC Manager, when you create a workload domain or add a vSphere cluster to a workload domain, the vSAN Performance Service is not enabled for vSAN clusters. When CEIP is enabled, data from the vSAN Performance Service is provided to VMware and this data is used to aid VMware Support with troubleshooting and for products such as VMware Skyline, a proactive cloud monitoring service. See Customer Experience Improvement Program for more information on the data collected by CEIP.

  Enable CEIP in SDDC Manager. See the VMware Cloud Foundation Documentation. After CEIP is enabled, a scheduled task that enables the vSAN Performance Service on existing clusters in workload domains runs every three hours. The service is also enabled for new workload domains and clusters. To enable the vSAN Performance Service immediately, see the VMware vSphere Documentation.

- **Creation or expansion of a vSAN cluster with more than 32 hosts fails**

  By default, a vSAN cluster can grow up to 32 hosts. With large cluster support enabled, a vSAN cluster can grow up to a maximum of 64 hosts. However, even with large cluster support enabled, a creation or expansion task can fail on the sub-task **Enable vSAN on vSphere Cluster**.

  Workaround:

  a Enable Large Cluster Support for the vSAN cluster in the vSphere Client. If it is already enabled skip to step 2.

    1 Select the vSAN cluster in the vSphere Client.

    2 Select **Configure > vSAN > Advanced Options**.

    3 Enable Large Cluster Support.

    4 Click **Apply**.

    5 Click **Yes**.

  b Run a vSAN health check to see which hosts require rebooting.

  c Put the hosts into Maintenance Mode and reboot the hosts.

  For more information about large cluster support, see https://kb.vmware.com/kb/2110081.

- **Removing a host from a cluster, deleting a cluster from a workload domain, or deleting a workload domain fails if Service VMs (SVMs) are present**

  If you deployed an endpoint protection service (such as guest introspection) to a cluster through NSX-T Data Center, then removing a host from the cluster, deleting the cluster, or deleting the workload domain containing the cluster will fail on the subtask `Enter Maintenance Mode on ESXi Hosts`.

  Workaround:

  - For host removal: Delete the Service VM from the host and retry the operation.
  - For cluster deletion: Delete the service deployment for the cluster and retry the operation.
  - For workload domain deleting: Delete the service deployment for all clusters in the workload domain and retry the operation.

- **vCenter Server overwrites the NFS datastore name when adding a cluster to a VI workload domain**

  If you add an NFS datastore with the same NFS server IP address, but a different NFS datastore name, as an NFS datastore that already exists in the workload domain, then vCenter Server applies the existing datastore name to the new datastore.

  Workaround: If you want to add an NFS datastore with a different datastore name, then it must use a different NFS server IP address.

- **Deleting a cluster that was renamed in the vSphere Client does not delete the cluster's transport node profile or uplink profile**

  When you use the vSphere Client to rename a cluster and then delete that cluster from SDDC Manager, the transport node profile and uplink profile associated with the cluster are not removed.

  Workaround: Manually delete the transport node profile and uplink profile from NSX Manager and try deleting the cluster again.

  a Log in to the NSX Manager UI.

  b Identify and delete the uplink profile associated with the cluster's old name.

    1 Navigate to **System > Fabric > Profiles > Uplink Profiles** and identify the uplink profile for the deleted cluster.

      Uplink profile names follow the pattern: **<vcenter host name>-<old-cluster-name>**.

      For example, if the vCenter Server's FQDN is **vcenter-vsan.vrack.vsphere.local**, and the cluster's old name is **nsxt-datacenter**, then uplink profile name would be **vcenter-vsan-nsxt-cluster**.

    2 Select the uplink profile and click **Delete**.

  c Identify and delete the transport node profile associated with the cluster's old name.

    1 Navigate to **System > Fabric > Profiles > Transport Node Profiles** and identify the transport node profile for the deleted cluster. Transport node profile names follow the same pattern as uplink profile names.

    2 Select the transport node profile and click **Delete**.

  d In SDDC Manager, try deleting the cluster again.

- **Add host may fail for cluster using vLCM images**

  Add host for cluster using vLCM images may fail with the following error: `Unable to create transport node collection. Transport Node Collection is failed with state FAILED_TO_REALIZEWorkflow`

  Workaround: Re-try the add host workflow.

Multi-Instance Management Known Issues

- **Federation creation information not displayed if you leave the Multi-Instance Management Dashboard**

  Federation creation progress is displayed on the Multi-Instance Management Dashboard. If you navigate to another screen and then return to the Multi-Instance Management Dashboard, progress messages are not displayed. Instead, an empty map with no Cloud Foundation instances are displayed until the federation is created.

Stay on the Multi-Instance Dashboard till the task is complete. If you have navigated away, wait for around 20 minutes and then return to the dashboard by which time the operation should have completed.

- **Multi-Instance Management Dashboard operation fails**

  After a controller joins or leaves a federation, Kafka is restarted on all controllers in the federation. It can take up to 20 minutes for the federation to stabilize. Any operations performed on the dashboard during this time may fail.

  Re-try the operation.

- **Join operation fails**

  A join operation may fail if a controller SDDC Manager has a public certificate with a depth greater than one (that is, it has intermediate certificates).

  Workaround: Trust the intermediate certificate of the controller SDDC Manager. See KB 80986.

API Known Issues

- **Stretch cluster operation fails**

  If the cluster that you are stretching does not include a powered-on VM with an operating system installed, the operation fails at the "Validate Cluster for Zero VMs" task.

  Make sure the cluster has a powered-on VM with an operating system installed before stretching the cluster.

vRealize Suite Known Issues

- **vRealize Operations Manager: VMware Security Advisory VMSA-2021-0018**

  VMSA-2021-0018 describes security vulnerabilities that affect VMware Cloud Foundation.

  - The vRealize Operations Manager API contains an arbitrary file read vulnerability. A malicious actor with administrative access to vRealize Operations Manager API can read any arbitrary file on server leading to information disclosure.The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned identifier CVE-2021-22022 to this issue.
  - The vRealize Operations Manager API has insecure object reference vulnerability. A malicious actor with administrative access to vRealize Operations Manager API may be able to modify other users information leading to an account takeover.The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned identifier CVE-2021-22023 to this issue.

- The vRealize Operations Manager API contains an arbitrary log-file read vulnerability. An unauthenticated malicious actor with network access to the vRealize Operations Manager API can read any log file resulting in sensitive information disclosure. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned identifier CVE-2021-22024 to this issue.
- The vRealize Operations Manager API contains a broken access control vulnerability leading to unauthenticated API access. An unauthenticated malicious actor with network access to the vRealize Operations Manager API can add new nodes to existing vROps cluster. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned identifier CVE-2021-22025 to this issue.
- The vRealize Operations Manager API contains a Server Side Request Forgery in multiple end points. An unauthenticated malicious actor with network access to the vRealize Operations Manager API can perform a Server Side Request Forgery attack leading to information disclosure. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned identifiers CVE-2021-22026 and CVE-2021-22027 to this issue.

Workaround: See KB 85452 for information about applying vRealize Operations Security Patches that resolve the issues.

- **vRealize Log Insight: VMSA-2021-0019**

  VMSA-2021-0019 describes security vulnerabilities that affect VMware Cloud Foundation.

  VMware vRealize Log Insight contains a Cross Site Scripting (XSS) vulnerability due to improper user input validation. An attacker with user privileges may be able to inject a malicious payload via the Log Insight UI which would be executed when the victim accesses the shared dashboard link. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned identifier CVE-2021-22021 to this issue.

  Workaround: See KB 85405 for information about applying a vRealize Log Insight Security Patch that resolves the issue.

- **When you deploy a vRealize Suite product in vRealize Suite Lifecycle Manager, some of the infrastructure details may not be loaded**

  When you deploy a vRealize Suite product in vRealize Suite Lifecycle Manager, the Infrastructure details may not get populated. This may indicate a failed vCenter data collection request in vRealize Suite Lifecycle Manager, which prevents vRealize Suite Lifecycle Manager from validating the current state of vCenter Server inventory.

  Workaround: Retry the failed vCenter data collection request until it successfully passes and then try to complete the vRealize Suite product deployment again.

  a In vRealize Suite Lifecycle Manager, click **Lifecycle Operations > Datacenters**.

  b Click the refresh icon for the management vCenter Server to refresh data collection.

  c Click **Requests** and wait for request to be successful.

  d Retry the product deployment.

- **Updating the DNS or NTP server configuration does not apply the update to vRealize Automation**

  Using the Cloud Foundation API to update the DNS or NTP servers does not apply the update to vRealize Automation due to a bug in vRealize Suite Lifecycle Manager.

  Workaround: Manually update the DNS or NTP server(s) for vRealize Automation.

  Update the DNS server(s) for vRealize Automation

  a SSH to the first vRealize Automation node using root credentials.

  b Delete the current DNS server using the following command:

  ```
  sed '/nameserver.*/d' -i /etc/resolv.conf
  ```

  c Add the new DNS server IP with following command:

  ```
  echo nameserver [DNS server IP] >> /etc/resolv.conf
  ```

  d Repeat this command if there are multiple DNS servers.

  e Validate the update with the following command:

  ```
  cat /etc/resolv.conf
  ```

  f Repeat these steps for each vRealize Automation node.

  Update the NTP server(s) for vRealize Automation

  a SSH to the first vRealize Automation node using root credentials.

  b Run the following command to specify the new NTP server:

  ```
  vracli ntp systemd --set [NTP server IP]
  ```

  To add multiple NTP servers:

  ```
  vracli ntp systemd --set [NTP server 1 IP,NTP server 2 IP]
  ```

  c Validate the update with the following command:

  ```
  vracli ntp show-confi
  ```

  d Apply the update to all vRealize Automation nodes with the following command:

  ```
  vracli ntp apply
  ```

  e Validate the update by running the following command on each vRealize Automation node:

  ```
  vracli ntp show-config
  ```

- **Connecting vRealize Log Insight to a workload domain fails at the "Enable Log Collection for vSphere" step**

When you connect vRealize Log Insight to a workload domain, it fails at the `Enable Log Collection for vSphere` step. Expanding the connect workflow task in the SDDC Manager Recent Tasks widget displays the following errors:

```
Cannot configure vCenter in vRealize Log Insight

Failed post request in vRLI
```

Workaround: Reboot the vRLI VMs.

a Log in to the management vCenter Server.

b Click **Virtual Machines** in the VMware Host Client inventory.

c Select a vRealize Log Insight VM.

d Right-click the VM and select **Guest OS** -> **Restart**.

e After the VM has been powered on, repeat steps 3 and 4 for each vRealize Log Insight VM.

f Validate that the vRealize Log Insight UI is accessible.

g In the SDDC Manager Recent Tasks panel, select the failed task and click **Restart Task**.

- **vRealize Suite product deployment fails with error "Failed to get Environment ID by given Host Name"**

  When you deploy multiple vRealize Suite products simultaneously in vRealize Suite Lifecycle Manager, product registration with SDDC Manager may fail.

  Workaround: Deploy vRealize Suite products one at a time.

- **Connecting vRealize Operations Manager to a workload domain fails at the "Create vCenter Server Adapter in vRealize Operations Manager for the Workload Domain" step**

  When you connect vRealize Operations Manager to a workload domain, it fails at the `Create vCenter Server Adapter in vRealize Operations Manager for the Workload Domain` step with a message similar to `Failed to configure vCenter <vcenter-hostname> in vROps <vrops-hostname>, because Failed to manage vROps adapter`. This issue can occur when the vRealize Operations cluster is offline.

  Workaround: Make sure that the vRealize Operations cluster is online.

  a Log in to the vRealize Operations Manager administration interface.

  b Click **Administration > Cluster Management** and check the cluster status.

  c If the vRealize Operations cluster is offline, bring the cluster online.

  d When the cluster status displays as online, retry connecting vRealize Operations Manager to a workload domain.