

VMware Cloud Foundation 4.1 Release Notes

VMware Cloud Foundation 4.1

VMware Cloud Foundation 4.1.0.1

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

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

Contents

1	What's New	4
2	VMware Cloud Foundation Bill of Materials (BOM)	6
3	VMware Software Edition License Information	8
4	Supported Hardware	9
5	Documentation	10
6	Browser Compatibility and Screen Resolutions	11
7	Installation and Upgrade Information	12
8	VMware Cloud Foundation 4.1.0.1 Release Information	13
9	Resolved Issues	15
10	Known Issues	16
	VMware Cloud Foundation Known Issues	16
	Upgrade Known Issues	18
	Bring-up Known Issues	21
	SDDC Manager Known Issues	21
	Workload Domain Known Issues	23
	Multi-Instance Management Known Issues	27
	API Known Issues	28
	Networking Known Issues	28
	vRealize Suite Known Issues	29

What's New

1

These releases have been determined to be impacted by CVE-2020-4006. Fixes and Workarounds are available to address this vulnerability. For more information, see [VMSA-2020-0027](#).

The VMware Cloud Foundation (VCF) 4.1 release includes the following:

- **vVols as Principal Storage in Workload Domains:** VMware Cloud Foundation now supports vVols as principal storage, providing a common storage management framework for external storage and automation for pre-defined storage, including volume management and provisioning.
- **Remote Clusters:** Extends VMware Cloud Foundation capabilities to the ROBO and Edge sites with VMware Cloud Foundation Remote Clusters. Now customers can enjoy the same consistent cloud operations in their core data center and edge/ ROBO sites.
- **Read-only Access and Local Accounts:** Administrators can create VIEWER users that have read-only access to VMware Cloud Foundation. They can also create a local account for use in break-glass scenarios where a remote identity provider is unreachable.
- **ESXi Parallel Upgrades:** Enables you to update the ESXi software on multiple clusters in the management domain or a workload domain in parallel. Parallel upgrades reduce the overall time required to upgrade your environment.
- **NSX-T Data Center Parallel Upgrades:** Enables you to upgrade all Edge clusters in parallel, and then all host clusters in parallel. Parallel upgrades reduce the overall time required to upgrade your environment.
- **Support for ESXi hosts with external CA-signed certificates:** VMware Cloud Foundation supports APIs to perform bring-up of hosts with certificates generated by an external Certificate Authority (CA).
- **vRealize Suite Lifecycle Manager in VMware Cloud Foundation mode:** VMware Cloud Foundation 4.1 introduces an improved integration with vRealize Suite Lifecycle Manager. When vRealize Suite Lifecycle Manager in VMware Cloud Foundation mode is enabled, the behavior of vRealize Suite Lifecycle Manager is aligned with the VMware Cloud Foundation architecture.

- **vSphere Cluster Services (vCLS) Support:** vCLS is a new capability introduced in the vSphere 7 Update 1 release. vCLS ensures that if vCenter Server becomes unavailable, cluster services remain available to maintain the resources and health of the workloads that run in the clusters.
- **Support for Renaming VMware Cloud Foundation Objects:** You can rename workload domains, network pools, and compute clusters after you have deployed them. This allows the flexibility of naming these Cloud Foundation objects to align with company policies.
- **VMware Skyline Support for VMware Cloud Foundation:** VMware Skyline brings proactive intelligence to VMware Cloud Foundation by identifying management and workload domains, and proactively surfacing VMware Cloud Foundation solution findings.
- **Backup Enhancements:** SDDC Manager backup and recovery workflows and APIs have been improved to add new capabilities including, backup management, backup scheduling, retention policy, on-demand backup, and automatic retries on failure. The enhancements also include Public APIs for 3rd party ecosystem and certified backup solutions from Dell PowerProtect and Cohesity.
- **Lifecycle Management Enhancements:** VMware Cloud Foundation allows skipping versions during upgrade to minimize the number of upgrades applied and time consumed in upgrading. Skip-level upgrade is managed using SDDC Manager and the public API.
- **Improved pNIC/vDS support:** VI Workload domains can have hosts with multiple pNICs and vSphere Distributed Switches (vDS) that can scale up-to the vSphere maximums supported in the vSphere version included in the BOM.
- **Support for XLarge form factor for Edge nodes:** You can now use SDDC Manager to create an edge cluster with the XLarge form factor for edge nodes in the Management and VI workload domains.
- **Localization:** SDDC Manager includes localization support for the following languages - German, Japanese, Chinese, French and Spanish. Customers can navigate the SDDC Manager UI in those languages.
- **Inclusive terminology:** As part of a [company-wide effort](#) to remove instances of non-inclusive language in our products, the VMware Cloud Foundation team has made changes to some of the terms used in the product UI and documentation.
- **New License for vSphere with Tanzu:** vSphere with Tanzu has its own license key, separate from vSphere 7.0. This is a subscription-based license with a term limit.
- **Start up and shut down order guidance:** [Start up and shut down order guidance](#) for VMware Cloud Foundation is now available, enabling you to gracefully shut down and start up the SDDC components in a prescriptive order.
- **Voluntary Product Accessibility Template (VPAT) report:** The VPAT evaluates compliance with accessibility guidelines as put forward by the US government (under Section 508) and the EU government (under EN 301 549). See <https://www.vmware.com/help/accessibility.html>.
- **BOM Updates:** Updated Bill of Materials with new product versions.

VMware Cloud Foundation Bill of Materials (BOM)

2

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

Software Component	Version	Date	Build Number
Cloud Builder VM	4.1.0.0	06 OCT 2020	16961769
SDDC Manager	4.1.0.0	06 OCT 2020	16961769
VMware vCenter Server Appliance	7.0 Update 1	06 OCT 2020	16860138
VMware ESXi	7.0 Update 1	06 OCT 2020	16850804
VMware vSAN	7.0 Update 1	06 OCT 2020	16850804
VMware NSX-T Data Center	3.0.2	17 SEP 2020	16887200
VMware vRealize Suite Lifecycle Manager	8.1 Patch 1	25 AUG 2020	16776528
Workspace ONE Access	3.3.2	14 APR 2020	15951611
vRealize Automation	8.1 Patch 2	28 JUL 2020	16633378
vRealize Log Insight	8.1.1	28 MAY 2020	16281169
vRealize Log Insight Content Pack for NSX-T	3.9.0	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	n/a	n/a
vRealize Log Insight Content Pack for VMware Identity Manager	2.0	n/a	n/a

Software Component	Version	Date	Build Number
vRealize Operations Manager	8.1.1	09 JUL 2020	16522874
vRealize Operations Management Pack for VMware Identity Manager	1.1	n/a	n/a

- 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 6.1 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

3

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

4

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

5

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

6

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

7

You can install Cloud Foundation 4.1 as a new release or upgrade to Cloud Foundation 4.1 with a sequential or skip-level upgrade. For a sequential upgrade to Cloud Foundation 4.1, your environment must be at Cloud Foundation 4.0.1.1. You can perform a skip-level upgrade from an earlier version of Cloud Foundation.

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 4.1

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

Upgrading to VMware Cloud Foundation 4.1

You can upgrade to Cloud Foundation 4.1 from 4.0.1.1. You can also use the skip-level upgrade tool to upgrade to VMware Cloud Foundation 4.1 from versions earlier than 4.0.1.1. For more information see [VMware Cloud Foundation Lifecycle Management](#).

VMware Cloud Foundation 4.1.0.1 Release Information



VMware Cloud Foundation 4.1.0.1 includes bug and security fixes.

You can upgrade to Cloud Foundation 4.1.0.1 from a 4.1 deployment. For upgrade information, refer to the VMware Cloud Foundation Lifecycle Management Guide.

Cloud Foundation 4.1.0.1 contains the following BOM updates:

Software Component	Version	Date	Build Number
SDDC Manager	4.1.0.1	24 NOV 2020	17206953
VMware ESXi	ESXi 7.0 Update 1b	19 NOV 2020	17168206

SDDC Manager 4.1.0.1 addresses the following issue:

- SDDC Manager 4.1.0.1 contains security fixes for Photon OS packages [PHSA-2020-3.0-0140](#) to [PHSA-2020-3.0-0162](#) published here: <https://github.com/vmware/photon/wiki/Security-Advisories-3>.

ESXi 7.0 Update 1b addresses the following issues:

- OpenSLP as used in ESXi has a use-after-free issue. This issue might allow a malicious actor with network access to port 427 on an ESXi host to trigger a use-after-free in the OpenSLP service resulting in remote code execution. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned the identifier [CVE-2020-3992](#) to this issue. For more information, see VMware Security Advisory [VMSA-2020-0023.1](#).
- VMware ESXi contains a use-after-free vulnerability in the XHCI USB controller. A malicious actor with local administrative privileges on a virtual machine might exploit this issue to execute code as the virtual machine's VMX process running on the host. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned the identifier [CVE-2020-4004](#) to this issue. For more information, see [VMSA-2020-0026](#).

- VMware ESXi contains a privilege-escalation vulnerability that exists in the way certain system calls are being managed. A malicious actor with privileges within the VMX process only, might escalate their privileges on the affected system. Successful exploitation of this issue is only possible when chained with another vulnerability. The Common Vulnerabilities and Exposures project (cve.mitre.org) has assigned the identifier [CVE-2020-4005](#) to this issue. For more information, see [VMSA-2020-0026](#).

Resolved Issues

9

The following issues are resolved in this release:

- Credential logging vulnerability as described in [VMSA-2022-0003](#). See [KB 87050](#) for more information.
- Validation fails when using the VMware Cloud Foundation API to commission an ESXi host with VMFS on FC storage.
- Adding an NSX-T Edge cluster fails if the SDDC Manager reboots before the task completes.
- vSAN File Services cannot be enabled on vLCM-enabled workload domains.
- Unable to remove hosts from a cluster that was unsuccessfully stretched.
- Unstretch cluster operation fails at task "Get Data from Inventory".
- After you create an NSX-T Edge cluster, SDDC Manager does not support expanding or shrinking it by adding or deleting Edge nodes. While this is still unsupported in SDDC Manager, you can contact VMware Support if you need to expand or shrink an NSX-T Edge cluster.

This chapter includes the following topics:

- [VMware Cloud Foundation Known Issues](#)
- [Upgrade Known Issues](#)
- [Bring-up Known Issues](#)
- [SDDC Manager Known Issues](#)
- [Workload Domain Known Issues](#)
- [Multi-Instance Management Known Issues](#)
- [API Known Issues](#)
- [Networking Known Issues](#)
- [vRealize Suite Known Issues](#)

VMware Cloud Foundation Known Issues

- **VMware Cloud Foundation and vRealize Automation multi-tenancy**

VMware Cloud Foundation 4.1 does not support multi-tenancy for vRealize Automation and Workspace ONE Access when using the vRealize Suite Lifecycle Manager deployed by SDDC Manager.

Workaround: None.

- **VMware vSAN HCI Mesh and VMware Cloud Foundation**

VMware Cloud Foundation 4.1 does not support VMware vSAN HCI Mesh.

Workaround: None

- **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.

■ Stretched clusters and Workload Management

You cannot stretch a cluster on which Workload Management (vSphere with Tanzu) is deployed.

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

■ The Bill of Materials (BOM) does not include the latest version of the vRealize Log Insight Content Pack for NSX-T

The BOM contains vRealize Log Insight Content Pack for NSX-T 3.9.0, even though version 3.9.1 is available.

Workaround: Install vRealize Log Insight Content Pack for NSX-T 3.9.1.

- a In a Web browser, log in to the vRealize Log Insight user interface.
- b Click the configuration drop-down menu icon and select **Content Packs**.
- c In the Navigator, under Content Pack Marketplace, click **Updates**.
- d On the Log Insight Content Pack Marketplace page, select the vRealize Log Insight Content Pack for NSX-T and click **Update**.
- e On the Update All Content Packs dialog, click **Update**.
- f Click on "NSX-T – VMware" under Installed Content Packs to verify that the version number is updated.

NOTE: The Content Pack Marketplace always includes the latest available version of Content Packs. You can install the latest version of the vRealize Log Insight Content Pack for NSX-T, even if it is not 3.9.1

■ The bundle transfer utility included as part of the SDDC Manager VM does not work as expected

If you do not have internet connectivity in your VMware Cloud Foundation system, you can use the bundle transfer utility to manually download bundles. The SDDC Manager VM includes version 1696170 of the utility. This version of the utility fails to download/list bundles when you specify the product version (`-p` or `--productVersion`).

Workaround: Download version 17209083 of the Bundle Transfer Utility & Skip Level Upgrade Tool from [MyVMware](#).

Upgrade 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.

■ Applying the configuration drift upgrade bundle fails

If the VMware Cloud Foundation environment that you are upgrading includes any failed or partially-deployed workload domains, applying the configuration drift bundle will fail.

Workaround:

- Perform an upgrade [precheck](#) and resolve any issues.
- Delete any partially-deployed workload domains.
- Apply the configuration drift bundle again.

■ 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

■ Skip-level upgrade and restarting the LCM service fail

If you applied a hot patch provided by VMware Engineering to Cloud Foundation 4.0, 4.0.0.1, or 4.0.1 then the following tasks may fail due to a version aliasing issue:

- Skip-level upgrade to VMware Cloud Foundation 4.1
- Restarting the LCM service (`systemctl restart lcm`)

Workaround: Contact VMware Support to resolve the version aliasing issue.

- **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 not support scratch partitions. See VMware KB article [2074026](#).

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

- Log in to SDDC Manager as `vcf` using a Secure Shell (SSH) client.
- Open the `application-prod.properties` file for editing.

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

- Add the following property and save the file.

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

- Restart the life cycle management service.

```
systemctl restart lcm
```

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

- **vRealize Operations Manager upgrade fails on the step `VREALIZE_UPGRADE_PREPARE_BACKUP` with the error: Waiting for vRealize Operations cluster to change state timed out**

When upgrading vRealize Operations Manager, SDDC Manager takes the vRealize Operations Manager cluster offline and takes snapshots of the vRealize Operations Manager virtual machines. In some circumstances, taking the cluster offline takes a long time and the operation times out.

Workaround: Take the vRealize Operations Manager cluster back online and retry the upgrade.

- Log in to the vRealize Operations Manager Administration UI (https://<vrops_ip>/admin) using the admin credentials.
- If the cluster status is offline, in the Cluster Status section click **Take Cluster Online**. Wait for the cluster to initialize and be marked as green.

- c In the SDDC Manager UI, the option to retry vRealize Operations Manager upgrade should be available. Retry the upgrade.

If the upgrade continues to fail, take the snapshots manually and retry the upgrade. Since the snapshots already exist, SDDC Manager will skip that step and proceed with the upgrade.

- a Log in to the vRealize Operations Manager Administration UI (https://<vrops_ip>/admin) using the admin credentials.
- b Ensure that the vRealize Operations Manager Cluster Status is offline. If it is online, click **Take Cluster Offline** in the Cluster Status section. Wait for the cluster to be marked as offline.
- c Log in to the management domain vCenter Server using the vSphere Client. Navigate to the vRealize Operations Manager virtual machines and create a snapshot for each virtual machine in the vRealize Operations Manager cluster. Use the following prefix "vROPS_LCM_UPGRADE_MANUAL_BACKUP" for the snapshots. Please note that the prefix should match the letter casing.
- d After the snapshots are done, log in to the vRealize Operations Manager UI and take cluster online. Wait for the cluster initialization.
- e In the SDDC Manager UI, the option to retry vRealize Operations Manager upgrade should be available. Retry the upgrade.

- **Cluster level upgrade is not available if the workload domain has a faulty cluster**

This issue occurs if any host or cluster in the workload domain is in an error state.

Workaround: Remove the faulty host or cluster from the workload domain. The cluster level upgrade option is then available for the workload domain.

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.

Workaround: 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

- **When you replace the certificates for NSX Manager in SDDC Manager the NSX Container Plug-in (NCP) crashes**

You will not be able to deploy new vSphere pods, load balancers, or other NSX-T resources until you restart the workload management service.

Workaround:

- a SSH into the vCenter Server appliance.
- b Run the following command:

```
vmon-cli -r wcp
```

■ **Certificate installation fails for NSX Manager**

You can't use CA-signed certificates that have LDAP-based CDPs (CRL Distribution Point).

Workaround: See KB article [78794](#).

■ **When you add a cluster to a workload domain that has a separate vSphere Distributed Switch (vDS) for overlay traffic, it may not have the correct mapping between the uplinks and pNICs**

If you have a workload domain that includes ESXi hosts with more than two pNICs and has multiple vSphere Distributed Switches, the uplinks for the overlay vDS may not map to the correct pNICs. In addition, if you create an Edge cluster for the workload domain before correcting the mapping, then BGP peering will fail.

Workaround: Use the vSphere Client to update the uplink names to match the actual uplinks on your ESXi hosts.

- a In the vSphere Client Home page, click **Networking** and select the distributed switch.
- b On the **Configure** tab, expand **Settings** and select **Properties**.
- c Click **Edit**.
- d Click **General**, then click **Edit uplink names** to change the names of the uplinks.

If the uplinks were mapped incorrectly and you created an Edge cluster, BGP peering will fail. Perform the following steps to resolve the issue:

- a In the vSphere Client Home page, click **Networking** and select the distributed switch.
- b On the **Configure** tab, expand **Settings** and select **Properties**.
- c Click **Edit**.
- d Click **General**, then click **Edit uplink names** to change the names of the uplinks so the pNICs are correctly mapped.

■ **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.

- **The Password Management page in SDDC Manager shows the incorrect account type for the PSC SSO administrator**

On the Password Management page in SDDC Manager, the account type for the PSC SSO administrator account displays as USER. It is actually a SYSTEM account. This is a UI-only issue and there is no functional impact.

Workaround: None.

- **The Workload Management wizard redirects you to the incorrect vCenter Server in the vSphere Client**

When you click the **Complete in vSphere** link in the Workload Management wizard, SDDC Manager always redirects you to the management vCenter Server in the vSphere Client. If you are deploying Workload Management to a VI workload domain, this is not the correct vCenter Server from which to complete the deployment.

Workaround: In the vSphere Client, navigate to the vCenter Server for the VI workload domain to complete Workload Management deployment.

- **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.

Workload Domain Known Issues

- **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 vmkO 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.

- **VMware Cloud Foundation does not support Service VMs (SVMs) on vLCM-enabled workload domains**

You cannot deploy a Service VM to an NSX Manager that is associated with a workload domain that is using vSphere Lifecycle Manager (vLCM).

Workaround: None.

- **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.

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

- **Unassigned host upgrade is not supported**

When performing an upgrade using the VMware Cloud Foundation API, the upgradeSpec does not support the resourceType "UNASSIGNED_HOST". The API Reference Guide includes examples that use "UNASSIGNED_HOST", but it is not supported.

Workaround: None.

- **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.

Networking Known Issues

- **VMware Cloud Foundation does not enable StandBy Relocation on Tier-1 gateways**

If you create an NSX-T Edge cluster with more than 2 Edge nodes, you should enable StandBy Relocation. Standby relocation means that if the Edge node where the active or standby logical router is running fails, a new standby logical router is created on another Edge node to maintain high availability. If the Edge node that fails is running the active logical router, the original standby logical router becomes the active logical router and a new standby logical router is created. If the Edge node that fails is running the standby logical router, the new standby logical router replaces it.

Workaround: Use the NSX Manager UI to enable StandBy Relocation for any Tier-1 gateway that is part of an NSX-T Edge cluster with more than 2 Edge nodes.

- **An outage of a top of rack switch in the data center might cause lack of availability of segments and services that are provided by NSX-T Data Center**

During the failover of a top of rack switch, the TEP communication between the NSX-T components is disrupted causing some segments and services to become unavailable.

Workaround: To ensure that NSX-T Edge TEP communication fails over to the second top of rack switch in the management or workload domain, modify the teaming policy of the port groups for the uplink traffic of the NSX-T Edge nodes.

- a In a Web browser, log in to vCenter Server by using the vSphere Client.
- b In the **Networking** inventory, expand the tree and browse to the vSphere Distributed Switch for the management domain.
- c In the navigation pane, right-click the port group for the first uplink and select **Edit Settings**.
- d In the **Edit Settings** dialog box, select **Teaming and failover**.
- e Move the uplink from **Unused uplinks** to **Standby uplinks** and click **OK**.
- f Repeat Step 5 and Step 6 for the other port group for edge uplink traffic in the management domain.
- g Repeat the procedure for the port groups for edge uplink traffic in the workload domain.

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.

■ **vRealize Automation/vRealize Operations Manager load balancer configuration is lost as a result of deleting a failed environment in vRealize Suite Lifecycle Manager**

This issue can occur after the following sequence of events:

- a You deploy vRealize Automation or vRealize Operations Manager through vRealize Suite Lifecycle Manager, but the deployment fails.
- b You create a new environment and try the deployment again, and this time it succeeds.
- c In vRealize Suite Lifecycle Manager, you delete the failed environment from the first attempt after the second attempt succeeded.

In this situation, deleting the failed environment deletes the load balancer used by the successful environment. To avoid this issue, always delete a failed vRealize Automation/vRealize Operations Manager environment from vRealize Suite Lifecycle Manager before trying to deploy again in a new environment.

Workaround: Manually recreate the load balancer configuration on the standalone Tier-1 router used for vRealize Automation or vRealize Operations Manager load balancing.

- **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.

- **After you replace a vCenter Server certificate from SDDC Manager the the vRealize Log Insight integration with vSphere stops working**

If you have integrated vRealize Log Insight with vSphere and you replace the certificate for vCenter Server from SDDC Manager, the integration stops working, since the new certificate is not trusted.

Workaround: In vRealize Log Insight, accept the new certificate.

- a Log in to vRealize Log Insight.
- b Navigate to the **Administration** tab.
- c Under Integration, click **vSphere**.
- d In the vCenter Server table, locate the vCenter Server for which you replaced the certificate and click the exclamation mark (!) in the Collection Status column to accept the new certificate.

- **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:

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

To add multiple NTP servers:

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

- c Validate the update with the following command:

```
vraccli ntp show-config
```

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

```
vraccli ntp apply
```

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

```
vraccli ntp show-config
```

■ **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