

VMware Cloud Foundation 4.0.1 Release Notes

VMware Cloud Foundation 4.0.1

VMware Cloud Foundation 4.0.1.1

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)	5
3	VMware Software Edition License Information	6
4	Supported Hardware	7
5	Documentation	8
6	Browser Compatibility and Screen Resolutions	9
7	Installation and Upgrade Information	10
8	VMware Cloud Foundation 4.0.1.1 Release Information	11
9	Resolved Issues	13
10	Known Issues	14
	Upgrade Known Issues	14
	Bring-up Known Issues	15
	SDDC Manager Known Issues	15
	Workload Domain Known Issues	16
	Multi-Instance Management Known Issues	21
	API Known Issues	21
	Networking Known Issues	21
	vRealize Suite Known Issues	22

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) 4.0.1 release includes the following:

- **NSX-T cluster-level upgrade support:** Users can upgrade specific host clusters within a workload domain so that the upgrade can fit into their maintenance windows. NOTE: The NSX Manager cluster is not upgraded until *all* host clusters in the workload domain are upgraded. New features introduced in the upgrade are not configurable until the NSX Manager cluster is upgraded.
- **Multi-pNIC/multi-vDS during bring-up:** The deployment parameter workbook now provides three vSphere Distributed Switch (vDS) profiles that allow you to perform bring-up of hosts with two or four physical NICs (pNICs) and to create up to two vSphere Distributed Switches for isolating VMkernel traffic.
- **Kubernetes in the management domain:** vSphere with Kubernetes is now supported in the management domain. With VMware Cloud Foundation Workload Management, you can deploy vSphere with Kubernetes on the management domain default cluster with only 4 hosts.
- **VMware HCX Product Interoperability with vSphere 7.0 (Preview):** VMware HCX R142 enables users to adopt VMware Cloud Foundation 4.0.1 as the destination environment. See the [VMware HCX Release Notes](#) for details and limitations.
- **L3 Aware IP Addressing (API only):** VI workload domains now support the ability to use hosts from different L2 domains to create or expand clusters that use vSAN, NFS, or VMFS on FC as principal storage.
- **BOM Updates:** Updated Bill of Materials with new product versions.

VMware Cloud Foundation Bill of Materials (BOM)

2

The 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	4.0.1.0	25 JUN 2020	16428904
SDDC Manager	4.0.1.0	25 JUN 2020	16428904
VMware vCenter Server Appliance	7.0.0b	23 JUN 2020	16386292
VMware ESXi	7.0b	23 JUN 2020	16324942
VMware vSAN	7.0b	23 JUN 2020	16324942
VMware NSX-T Data Center	3.0.1	23 JUN 2020	16404613
VMware vRealize Suite Lifecycle Manager	8.1 Patch 1	21 MAY 2020	16256499

- Cloud Foundation supports, but does not automate, the deployment of VMware Horizon 7 version 7.12. You can deploy Horizon 7.12 on a workload domain using the Horizon 7.12 documentation.
- You can use vRealize Suite Lifecycle Manager to deploy vRealize Automation 8.1, vRealize Operations Manager 8.1, and vRealize Log Insight 8.1 using the VMware Validated Design 6.0 documentation.
- VMware Enterprise PKS is not supported with this release of Cloud Foundation.

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 vSAN Ready Nodes in Cloud Foundation, see [VMware Compatibility Guide \(VCG\) for vSAN](#) 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 4.0.1 as a new release or upgrade from Cloud Foundation 4.0.0.1.

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 4.0.0.1

Refer to the [VMware Cloud Foundation Deployment Guide](#) for information on deploying Cloud Foundation.

VMware Cloud Foundation 4.0.1.1 Release Information



VMware Cloud Foundation 4.0.1.1 was released on 06 AUG 2020. You can upgrade to Cloud Foundation 4.0.1.1 from a 4.0.1 deployment.

VMware Cloud Foundation 4.0.1.1 contains the following BOM updates:

Software Component	Version	Date	Build Number
SDDC Manager	4.0.1.1	06 AUG 2020	16660200
VMware vCenter Server Appliance	7.0.0c	30 JULY 2020	16620007

VMware vCenter Server Appliance 7.0.0c includes the following new features:

- Supervisor cluster: new version of Kubernetes, support for custom certificates and PNID changes
 - The Supervisor cluster now supports Kubernetes 1.18.2 (along with 1.16.7 and 1.17.4)
 - Replacing machine SSL certificates with custom certificates is now supported
 - vCenter PNID update is now supported when there are Supervisor clusters in the vCenter
- Tanzu Kubernetes Grid Service for vSphere: new features added for cluster scale-in, networking and storage
 - Cluster scale-in operation is now supported for Tanzu Kubernetes Grid service clusters
 - Ingress firewall rules are now enforced by default for all Tanzu Kubernetes Grid service clusters
 - New versions of Kubernetes shipping regularly asynchronously to vSphere patches, current versions are 1.16.8, 1.16.12, 1.17.7, 1.17.8
- Network service: new version of NCP
 - SessionAffinity is now supported for ClusterIP services
 - IngressClass, PathType, and Wildcard domain are supported for Ingress in Kubernetes 1.18
 - Client Auth is now supported in Ingress Controller
- Registry service: new version of Harbor
 - The Registry service now is upgraded to 1.10.3

For more information and instructions on how to upgrade, refer to the [Updating vSphere with Kubernetes Clusters](#) documentation.

VMware vCenter Server Appliance 7.0.0c resolves the following issue:

- Tanzu Kubernetes Grid Service cluster NTP sync issue

Resolved Issues

9

The following issues have been resolved:

- Credential logging vulnerability as described in [VMSA-2022-0003](#). See [KB 87050](#) for more information.
- Bring-up fails if Edge node FQDNs or the DNS Zone Name contain uppercase letters
- Cloud Builder appliance platform audit issues
- The platform audit connectivity checks for the uplink IP addresses for the NSX-T Edge nodes (for Application Virtual Networks) are executed with random IPs from the respective networks.
- Bring-up fails if you select "extra small" for the NSX-T Edge Node Appliance Size
- Edge cluster deployment task is stuck with a "Running" status
- Operations on NSX-T workload domains fails if their host FQDNs include uppercase letters
- The /v1/sso-domains APIs are not supported
- Add host operations using the Cloud Foundation API fail if no ESXi license is specified
- You cannot execute APIs that require scheduledTimeStamp as an input from the API Explorer
- Cloud Foundation supports only one NSX-T Edge Cluster residing on a vSphere Cluster
- Edge nodes are not configured with DNS or NTP settings after Edge cluster creation
- Adding hosts from different network pools to NSX-T workload domain clusters is only supported for hosts using vSAN storage
- Configuring a backup schedule for the SDDC Manager VM fails

Known Issues

10

This chapter includes the following topics:

- [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](#)

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.

- **NSX-T Data Center or ESXi upgrade fails for Workload Management workload domains with DaemonSets**

When DaemonSets are added to a supervisor cluster, the ESXi hosts in the workload domain cannot enter maintenance mode and upgrades fail.

Workaround: Remove the DaemonSets from the supervisor cluster before upgrading. After the upgrade is complete, you can redeploy the DaemonSets.

NOTE: This issue is resolved in Cloud Foundation 4.0.1.1.

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
```

- **Bring-up fails during the step Configure NSX-T Transport Node Action**

Bring-up fails and reports **Failed configuring transport node for ESXi <server hostname>..** When logged into NSX Manager, one or more hosts will show either a failure to configure a transport node or remain in an unconfigured state.

Workaround:

- If the NSX Manager reports that one or more hosts failed to configure a transport node, resolve the errors on the failed host(s) in NSX Manager and then retry bring-up.
- If the NSX Manager reports that one or more hosts do not have a transport node configured, retry bring-up.

SDDC Manager Known Issues

- **Adding an NSX-T Edge cluster fails if the SDDC Manager reboots before the task completes**

If the SDDC Manager VM reboots before the Edge cluster task completes, the task will fail and cannot be restarted until you restart to domain manager service.

Workaround:

- SSH to the SDDC Manager VM.
- Run the following command: `systemctl restart domainmanager.service`
- Verify that the domain manager service is running: `systemctl status domainmanager.service`
- Retry the Edge cluster deployment from the SDDC Manager UI.

- **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:

- 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.

- **Compose a server task fails**

For new installations of Cloud Foundation 4.0.1, composing a server (**Administration > Composable Infrastructure**) fails. This issue does not impact environments that were upgraded to Cloud Foundation 4.0.1.

Workaround: Upgrade to Cloud Foundation 4.0.1.1, which resolves this issue.

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.

■ **You are not able to add a cluster or a host to a NSX-T workload domain that has a dead host**

If one of the hosts of the workload domain goes dead and if you try to remove the host, the task fails. And then, that particular host is set to the deactive state without an option to forcefully remove it. In this condition, if you try to add a new cluster or add a host to the workload domain, the task runs for a long time and then fails eventually.

Workaround: Bring the dead host back to normal state, after which you would be able add a cluster and a 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](#).

■ **vSAN File Services cannot be enabled on vLCM-enabled workload domains**

In vSphere 7.0, vSphere Lifecycle Manager (vLCM) and vSAN File Services cannot be simultaneously be enabled on a vSAN cluster. See the [VMware vSphere 7.0 Release Notes](#) for more details on this limitation.

Workaround: None.

■ **Unable to remove hosts from a cluster that was unsuccessfully stretched**

After a failed attempt to stretch a cluster, removing a host from the cluster fails at the task **Enter Maintenance Mode on ESXi Hosts**.

Workaround: Log in to the vSphere Client and put the ESXi host into maintenance mode manually, then retry the remove hosts task in the SDDC Manager UI.

■ **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.

■ **Unstretch cluster operation fails at task Get Data from Inventory**

If the cluster contains hostnames that are uppercase, the unstretch operation may fail.

Workaround:

- a Log in to the SDDC Manager VM as **vcf**.
- b Enter **su** to switch to root user.
- c Enter the following commands to change hostnames to lowercase in the SDDC Manager inventory:

```
1  psql -h localhost -U postgres
2  \connect platform
3  update host set hostname = lower(hostname);
```

- d Start a new unstretch cluster 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.

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.

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.

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.

■ **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