



# Workload Optimization Manager 3.3.2 Release Notes

May 6, 2022

This document describes issues that are addressed in Workload Optimization Manager 3.3.2 – Release Date: May 6, 2022. Starting with the 3.0 version family, builds are cumulative. Please see the Workload Optimization Manager documentation for earlier versions of the Release Notes.

For any questions, please contact your support representative.

## Configuring Kubernetes Targets for Workload Optimization Manager

To set up a Kubernetes target for Workload Optimization Manager, you will deploy the Kubeturbo pod with specific configuration resources. These resources require your version of Workload Optimization Manager, mapped to a `TURBONOMIC_SERVER_VERSION`. Use the following table to map your version of Workload Optimization Manager:

Workload Optimization Manager Version:	TURBONOMIC_SERVER_VERSION number
3.3.2	8.5.2
3.3.1	8.5.1
3.3.0	8.5.0

For information about configuring the Kubeturbo pod, see the Kubeturbo GitHub repository, located at <https://github.com/turbonomic/kubeturbo>.

For information about Kubernetes targets and other targets, see the "Workload Optimization Manager Target Configuration Guide".

## What's New for Version 3.3.2

### Version 3.3.2

- **Newly Supported AWS Instances**  
This release introduces support for AWS EC2 [C6a instances](#).
- **New Filters for VMs**

You can now filter VMs by target name, type, or category when you use Search or create groups. Target name supports regular expressions. Target type includes individual hypervisors, cloud providers, container platforms, and APM vendors. Target category includes groups of related targets.

#### ■ **Constraint for Sidecar Container Specs**

A Kubernetes service might include [sidecar](#) Container Specs to provide additional services to a running pod, such as security or logging services. Sidecars injected at pod creation cannot be updated from the parent Workload Controller, causing a resize action to fail.

To prevent the execution of resize actions on injected sidecars, Workload Optimization Manager now adds them to a new group called "Injected Sidecars/All ContainerSpecs". This group applies a read-only policy that sets the action mode for resizes to *Recommend*. This means that you can only execute resizes outside of Workload Optimization Manager. The parent Workload Controller will continue to resize non-sidecar Container Specs as usual.

## Versioning Explanation

Workload Optimization Manager versioning uses V-R-M elements (Version, Release, Modification) in the version number to express the status of a given release, as follows:

Numbered Element	Example	Description
<b>V</b> - Version number	3.X.X	<ul style="list-style-type: none"> <li>Changes to platform architecture or significant changes to data models</li> </ul>
<b>R</b> - Release number	X.1.X	<ul style="list-style-type: none"> <li>Major feature changes</li> </ul>
<b>M</b> - Modification number is zero (0)	X.X.0	<ul style="list-style-type: none"> <li>A quarterly release</li> <li>All preview features from previous bi-weekly releases are now GA</li> <li>No new preview features in this release</li> </ul>
<b>M</b> - Modification number is greater than zero (1 or higher)	X.X.3	<ul style="list-style-type: none"> <li>A bi-weekly release</li> <li>Can include new preview features</li> <li>Includes fixed issues</li> </ul>

#### **NOTE:**

For API developers, the X.X.1 release can include final implementations of deprecated API features. These final implementations can make API changes that are not backward-compatible.

#### **NOTE:**

##### **Deprecation Notice for Tbmigrate, the Classic-To-XL Migration Tool**

Starting with the 3.3.0 release of Workload Optimization Manager, we will no longer support the Classic-To-XL Migration Tool. This tool provides a scripted interface that you can use to migrate from Classic installations (the 2.x version families) to XL installations (the 3.x version families).

You can still use this tool to migrate from Classic to the 3.2.x version family. We recommend that you migrate to version 3.2.0, then update to the latest version.

If you need to perform a migration after we release Workload Optimization Manager 3.3.0, you will be instructed to install and migrate to version 3.2.0, and then update that installation to the later version.

For more information, see "Appendix: Migrating Workload Optimization Manager from Classic to XL" in the *Workload Optimization Manager Installation Guide*.

## Configuration Requirements

For this release of Workload Optimization Manager, you should satisfy the following configuration requirements.

## Workload Optimization Manager Updates and Operator Version

Workload Optimization Manager deploys as a cloud-native application on a Kubernetes cluster. This cluster can be pre-configured on a VM that you deploy, or you can deploy Workload Optimization Manager to a Kubernetes cluster in your environment. In either case, Workload Optimization Manager uses an *Operator* to manage the application deployment.

For different versions of Workload Optimization Manager, we change the version of Operator you should use, as follows:

Product Version:	Operator Version:
3.5.2	42.9
3.2.6 - 3.5.1	42.7
3.2.4 - 3.2.5	42.6
3.2.3	42.5
3.2.1 - 3.2.2	42.4
3.1.5 - 3.2.0	42.3
3.1.4	42.2
3.1.2 - 3.1.3	42.1
3.1.1	42.0
3.1.0	8.2

When you update Workload Optimization Manager, you should be sure to always include the matching version of Operator in the update. When you perform an online or offline update according to the instructions in the Installation Guide, the procedure should automatically include the latest Operator.

If you have installed Workload Optimization Manager on a Kubernetes cluster, you might need to manually update the Operator version.

### OpenShift Installations of Workload Optimization Manager

For versions of OpenShift 4.x or later, if you manage your installation via OpenShift Operator Hub, navigate the OPC Console to **Installed Operators**. Select the version of the Workload Optimization Manager Operator that you want to use, and click **Update**. This should update the Operator, and you can now update Workload Optimization Manager to the matching version.

### Other Kubernetes Installations of Workload Optimization Manager

For installations on other supported Kubernetes platforms, you can update the Operator version in either one of two ways:

- Directly edit the running deployment of Workload Optimization Manager:
  1. Enter EDIT mode for your cluster.
 

```
kubectl edit deployment t8c-operator -n {YourNamespace}
```
  2. Edit the Operator image.
 

Find and edit image: {required value}
  3. Validate that the Operator pod is ready.
 

Execute the command: `kubectl get pods -n {YourNamespace}`, and verify that the pod is running and ready.
- Edit the Workload Optimization Manager deployment YAML file:
  1. Open the Operator deployment file for editing.
 

In the location where you store your manifests, open the file `operator.yaml`. This should be the file that you use to deploy the `t8c-operator` pod.
  2. Edit the Operator image.
 

Find and edit image: {required value}
  3. Apply the change to the operator.
 

```
kubectl apply -f operator.yaml
```
  4. Validate that the Operator pod is ready.

Execute the command: `kubectl get pods -n {YourNamespace}`, and verify that the pod is running and ready.

After you have updated the Operator version, and you have verified that the pod is running and ready, you can edit your Custom Resource declaration to update Workload Optimization Manager to the version that matches your Operator version.

For more information, contact your support representative.

## Supported MariaDB Version

For its default historical database, Workload Optimization Manager currently supports MariaDB version 10.5.13. This support includes comprehensive testing and quality control for Workload Optimization Manager usage of the historical database.

### IMPORTANT:

Do not update MariaDB to later versions (10.5.14 or 10.5.15). Workload Optimization Manager does not support operation with versions later than 10.5.13.

If you are running Workload Optimization Manager installed as an OVA, and using the database that is included in that OVA installation, then you must use version 10.5.13. For versions of Workload Optimization Manager that you installed as an OVA *before* version 3.2.6, it is likely that you need to update your MariaDB.

For information about updating your MariaDB instance, see "Verifying your MariaDB Version" in the latest version of the *Workload Optimization Manager Installation Guide*.

## SQL Modes for External Databases

If you deploy Workload Optimization Manager to work with an external database instead of the included historical database, then you must specify the correct SQL modes for the database. Configure the database to support:

```
{ {ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION} }
```

In particular, the SQL modes should *not* include `ONLY_FULL_GROUP_BY`, `NO_ZERO_IN_DATE`, or `NO_ZERO_DATE`.

## Transport Layer Security Requirements

By default, Workload Optimization Manager requires Transport Layer Security (TLS) version 1.2 to establish secure communications with targets. Most targets should have TLS 1.2 enabled. However, some targets might not have TLS enabled, or they might have enabled an earlier version. In that case, you will see handshake errors when Workload Optimization Manager tries to connect with the target service. When you go to the Target Configuration view, you will see a Validation Failed status for such targets.

In particular, we have found that NetApp filers often have TLS disabled by default, and that the latest version they support is TLS 1.0. If your NetApp target fails to validate, this could be the cause.

If target validation fails because of TLS support, you might see validation errors with the following strings:

- No appropriate protocol
  - To correct this error, ensure that you have enabled the latest version of TLS that your target technology supports. If this does not resolve the issue, please contact Technical Support.
- Certificates does not conform to algorithm constraints
  - To correct this error, refer to the documentation for your target technology (for example, refer to NetApp documentation) for instructions to generate a certification key with a length of 1024 or greater on your target server. If this does not resolve the issue, please contact Cisco Technical Support.

# Improvements

- **Improvement:**

- **For VDI environments, BusinessUser charts now show percentile data.**

- For Horizon VDI environments, to better understand the actions that analysis generates, the BusinessUser scope now includes the VCPU Percentile and Avg Utilization and VMEM Percentile and Avg Utilization charts.

- **Improvement:**

**You can now set Instance Store Aware Scaling for Azure VMs.**

To ensure cloud scaling actions respect instance store requirements, Workload Optimization Manager supports **Instance Store Aware Scaling** for VM policies. Starting with this release, you can use this setting for Azure workloads.

- **Improvement:**

**Logs now show permission errors for GCP billing account targets.**

When discovery of a GCP billing account target fails due to permission-related errors (for example, insufficient or invalid permissions), logs now show the errors to help you troubleshoot the issue.

- **Improvement:**

**Confirmation dialog box before you save changes to an automation policy.**

Starting with this version, the user interface displays a Confirmation dialog box before you save changes to an automation policy. This reminds you to consider the impact your changes might have on multiple entities in your environment.

## Fixed Issues

- **Fixed Issue:**

**Under some circumstances, an update to a newer version does not also update the version of Operator.**

Product updates should include updates to the Operator. If your update of Workload Optimization Manager does not include an update to Operator, then you can see an error in Target Status that says `Probe version (xxx) is not current; Please upgrade probe to version yyy.`

- **Fixed Issue:**

**For Azure, the user interface can show gaps in the billing data for RI utilization and coverage.**

Under some circumstances the Azure API that Workload Optimization Manager uses to get billing data can experience request interruptions. When this happens, charts can show a drop to zero for RI utilization and coverage.

- **Fixed Issue:**

**For Kubernetes, when ovn-kubernetes is used for an overlay network, a Pod Move action can fail to execute.**

For Kubernetes, when ovn-kubernetes is used for an overlay network, a Pod Move action can fail to execute.

- **Fixed Issue:**

**When you add a new host to your environment, under rare circumstances analysis can fail.**

When you add a new host to your environment, under rare circumstances analysis can fail with an error similar to:

```
FindInitialPlacement: Updating access commodity in historical economy cache encounter error...
```

- **Fixed Issue:**

**For Cloud-only environments, under some circumstances the user interface will not allow you to set scope in a chart.**

For Cloud-only environments, when you set scope for a chart the list of scopes to choose from can try to include on-prem scopes. However, if your environment has no on-prem entities, this list does not display any scopes, so you cannot choose a scope for the chart.

- **Fixed Issue:**

**Filters for search and groups based on target type do not support Azure Subscription or GCP Project.**

For VMs, when searching or creating a group based on target type, the user interface does not include options to filter by **Azure Subscription** or **GCP Project**. At this time you cannot create a filter on these criteria.

- **Fixed Issue:**

**For NetApp, discovery can fail overall if discovery encounters errors for a single entity.**

For NetApp environments, if discovery fails to get metrics for a single entity, then discovery can fail for the complete target.

- **Fixed Issue:**

**For Azure, storage resize actions can be incorrect.**

For Azure environments, if you have storage policies that require specific storage tiers, analysis can generate actions to incorrectly increase storage capacity. This can happen if a VM uses two volumes, and then you power down the VM.

- **Fixed Issue:**

**For Instana, the calculation for delayed discovery does not recognize a change to the discovery period for the probe.**

For Instana targets, you can change the discovery period for the probe. For example, if the default period is not sufficient for discovery to complete, then Workload Optimization Manager will show the target health as critical. To avoid this, you can lengthen the discovery period to give the probe time to complete discovery in your environment. However, the logic to show the probe health only recognizes the default discovery period.

- **Fixed Issue:**

**For the ServiceNow integration, action approval records do not always show the correct Category and Risk for the given actions.**

For the ServiceNow integration, action approval records do not always show the correct Category and Risk for the given actions.

- **Fixed Issue:**

**For the ServiceNow integration, action records do not show units of measure for memory values.**

For the ServiceNow integration, the records for actions show memory values, but they do not show the units of measure. They should display as MB or GB.

- **Fixed Issue:**

**For large vCenter environments, the component that runs the vCenter probe can run out of memory.**

For environments with many vCenter targets, the component that runs the vCenter probe can exceed the available count of threads. This causes the component to fail with an Out of Memory error.

- **Fixed Issue:**

**For VMs on Azure, analysis can recommend resizing to an SKU that the Guest OS cannot support.**

When scaling VMs in Azure, Workload Optimization Manager does not restrict the action to an SKU that the Guest OS can support. For example, it can recommend resizing to a VM that has more CPUs than the OS can support.

- **Fixed Issue:**

**For AWS, discovery does not recognize db.m6i or db.r6i workload tiers.**

For AWS environments, discovery does not recognize the db.m6i or db.r6i workload tiers. If you have AWS managed databases from those tiers in your environment, discovery will skip them, and you will see an error similar to:

```
Cannot find attributes for :db.m6i.8xlarge...
```

- **Fixed Issue:**

**If any platform components are down, you cannot set the product logging level.**

When you set the Workload Optimization Manager logging level, the user interface queries the logging levels currently set for all the platform components. If it cannot get data for one component, then it does not enable you to make the setting.

- **Fixed Issue:**

**For some VM Memory actions, the details can show incorrect amounts for reclaimed VCPU.**

For some VM Memory actions, the action details can incorrectly show large values for reclaimed VCPU. For example, for an action to remove a VCPU limit of 2000 MHZ, the action details can show that you reclaim 2000 VCPUs.

- **Fixed Issue:**

**The Volume Groups user interface does not filter multiple Resource Groups.**

When creating a group of volumes, if you set up a filter to include multiple Resource Groups in the search, the list of volumes is empty.

- **Fixed Issue:**

**Details for actions with very small savings or costs can show incorrect category (Potential Savings or Investment).**

Details for actions with very small savings or costs can show incorrect category (Potential Savings or Investment).

- **Fixed Issue:**

**When you display Price Overrides settings, performance is unacceptable.**

The user interface shows unacceptable performance when you navigate to **Price Adjustments** and then click to show **Price Overrides** for an entry.

- **Fixed Issue:**

**For Azure, the log is flooded with "missing database template" errors.**

In Azure environments, the log can be flooded with errors similar to:

```
[AzureDatabaseConverter] : Missing corresponding template for AzureRdbEntity
```

This error can occur when Workload Optimization Manager discovers unsupported database types in your environment.

## Known Issues

- **Known Issue:**

**For Azure environments with Ultra Disks, under rare circumstances discovery can fail to load data for specific entities.**

If the Azure environment returns null utilization data for a given Ultra disk, then discovery for the attached VM can fail with a `NoSuchElementException` error.

- **Known Issue:**

**For GCP, actions to scale workloads or to increase RI coverage show 0% for before and after.**

For Google Cloud Platform (GCP) environments, Workload Optimization Manager analysis does not currently use Committed Use Discount (CUD) data. As a result, action details show 0% for before and after costs or RI coverage. We will address this as our support for GCP develops.

- **Known Issue:**

**For Headroom plans, you can specify the desired state for overprovisioned commodities.**

Workload Optimization Manager runs nightly plans to calculate headroom in clusters – The number of VMs you can add to a cluster without exceeding limits for Memory, CPU, or storage. For example, if you want a desired state of 80% consumption, the plan will not add VMs that will exceed 80% utilization of resources in the cluster.

Note that reservations use these headroom calculations to determine whether Workload Optimization Manager can place the workloads in a reservation request.

By default, plans do not keep the utilization of overprovisioned resources within the desired state (80% in the above example). The plan calculates 100% utilization of overprovisioned resources. However, when placing reservations in the cluster, analysis could recommend provisioning new hosts as a way to keep the overprovisioned resources within the desired state.

We have introduced a setting you can make to enforce the desired state on overprovisioned resources. This can result in a lower calculated headroom. But when a reservation places and powers on a new VM, it is guaranteed to fit on the current infrastructure. Analysis will not see a need to provision a new host to support the VM's overprovisioned resources.

To turn on this feature, edit the topology-processor settings in the `cr.yaml` file:

1. Open an SSH terminal session on your Workload Optimization Manager instance.

Log in with the System Administrator that you set up when you installed Workload Optimization Manager.

2. Open the `cr.yaml` file for editing. For example:

```
vi /opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml
```

3. In the `spec/properties` section, find the entry for the `topology-processor` component.

4. Add the following properties to the component `spec`:

```
considerDesiredStateForProvisioningInClusterHeadroomPlan: true
considerUtilizationConstraintInClusterHeadroomPlan: true
```



5. Save and apply your changes to the platform.

After you save your changes, use `kubectl` to apply the changes:

```
kubectl apply -f /opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml
```

For assistance, contact your Support representative.

- **Known Issue:**

**For Kubernetes, in some environments analysis cannot execute Scale Node actions.**

For Kubernetes OCP 4.x and AKS environments, Workload Optimization Manager can generate and execute Scale Node actions. However, if the environment includes other K8s distributions that don't support execution of Scale Node actions (EKS, AKS, and OCP), then Workload Optimization Manager can disable execution of *all* Scale Node actions in the environment.

- **Known Issue:**

**For Kubernetes environments with nodes running Linux with cgroup v2 enabled, you must use Kubernetes version 1.23.2 or later.**

For Kubernetes environments with nodes running Linux with cgroup v2 enabled, an issue in earlier versions of Kubernetes prevents Workload Optimization Manager from collecting CPU utilization data for the affected nodes. To collect CPU utilization from the cgroup v2 nodes, you must run Kubernetes version 1.23.2 or later.

- **Known Issue:**

**For very large environments, under rare circumstances queries can lock the database and cause gaps in discovered data.**

For very large environments, it is possible for queries to lock the database, and cause gaps in discovered data. When this occurs you can see errors similar to the following:

```
- [RollupProcessor] : Error during rollup activity for table vm_stats_latest:
- [ComponentBasedTargetDumpingSettings] : Retaining 0 discovery dumps for target
```

If this occurs in your environment, contact your support representative for a workaround.

- **Known Issue:**

**Billed costs shown in Workload Optimization Manager and Google Cloud Platform (GCP) do not match due to time differences.**

Billed costs shown in Workload Optimization Manager charts and GCP billing reports do not match because Workload Optimization Manager uses UTC, while GCP uses local time. However, costs shown in both places are correct and reliable.

- **Known Issue:**

**For IBM FlashSystem, a known FlashSystem issue can cause some models to show intermittent target errors.**

For some models of IBM FlashSystem platforms, Workload Optimization Manager can show intermittent `Invalid Credentials` target errors. This is caused by a known FlashSystem REST service issue with platforms that have less than 64GB of memory.

If you experience this issue, use the following command to restart the FlashSystem REST service:

```
satask restartservice -service cfrest
```

Please contact your IBM FlashSystem Support representative for more information.

- **Known Issue:**

**For workloads with very large disks, Storage vMotion can time out.**

Storage vMotion actions for VMs with very large disks can time out. If you experience this, contact your support representative for help changing the timeout threshold.

- **Known Issue:**

**When updating your installation, under rare circumstances the Topology Processor component can fail to restart.**

When updating Workload Optimization Manager, under rare circumstances the `topology-processor` pod can fail to restart. The log posts an error with the statement:

```
AccessDeniedException: /home/turbonomic/data/kv
```

If you experience this problem, please contact your support representative for a workaround.

- **Known Issue:**

**For Azure, targets in the Australia Central region can show inconsistent pricing for VM license costs.**

For Azure environments running in the Australia Central region, the pricing for license costs that is reported to Workload Optimization Manager can be incorrect. As a result, the user interface can display incorrect license costs for Linux or Windows operating system licenses.

- **Known Issue:**

**For Azure and AWS, analysis does not consider the full cost for some workloads.**

For Azure environments, analysis considers the base OS cost, but does not consider additional costs for support or other add-on features that are bundled with the OS. The affected OS types are Ubuntu PRO, SUSE 24/7, and RHEL with HA.

For AWS environments, analysis does not consider AWS Marketplace costs.

- **Known Issue:**

**After updating the platform, Embedded Reports can fail to display.**

Under some circumstances, after you update Workload Optimization Manager to a new version, the Embedded Reports page can fail to display. The update appears to have completed successfully, but the Embedded Reports components appear to be running and ready. However, when you click on the Embedded Reports button, you get the following error:

```
failed to log in as user, specified in auth proxy header.
```

This can occur when the host VM takes a long time to completely start up all the platform components. If you encounter this problem, make sure all the components are running, and then execute the following command to restart the `grafana` pod:

```
kubectl delete pod -l app=grafana
```

For assistance, contact your support representative.

- **Known Issue:**

**For Citrix XenCenter environments with VMs that share storage, analysis can report excess capacity for thin-provisioned storage.**

For Citrix Hypervisor targets, analysis can calculate thin-provisioned capacity for storage. However, if multiple VMs share that storage, the calculated storage capacity can be doubled for each additional VM.

- **Known Issue:**

**For Fabric environments, when you scope the view to Datacenter the supply chain can fail to include all the related Host entities.**

For environments that include Fabric targets, under some circumstances a view that you scope to the Datacenter might not include all the related Host entities. This can happen for hosts that include the hyphen character ("-") in the Host name. For example, for Cisco UCS targets, the supply chain will not show UCS hosts (that use the hyphen character in their names) when you scope to the Datacenter entity.

- **Known Issue:**

**For Azure, under some circumstances a successful scaling action appears in the logs as failed.**

For Azure environments that include Availability Sets, under some circumstances a successful scaling action appears in the Workload Optimization Manager logs as failed.

- **Known Issue:**

**With Executed Actions charts, some data is missing for actions on entities that have been removed from the environment.**

When you view Executed Actions charts or export data from them, some data is missing for actions on entities that have been removed from the environment. For example, assume an action was executed on a storage volume, and that volume has later been removed from the environment. In that case, the exported data for that action will not include values that describe the removed volume.

- **Known Issue:**

**For Azure, if you connect the target through a proxy, then the target does not discover unattached storage volumes.**

For Azure environments, if you connect to the Azure target through a proxy, then Workload Optimization Manager does not discover unattached volumes.

- **Known Issue:**

**The Onboarding wizards can sometimes fail to close.**

When you first install Workload Optimization Manager, the user interface displays onboarding wizards to walk you through setting up your license, and configuring your first target. Under some circumstances, the button to end the wizard's workflow does not close the wizard. That can result in blocking you from continuing your Workload Optimization Manager session.

If the onboarding wizard does not close when you click **End Setup**, refresh the browser. That should close the wizard and leave you on the last user interface page that you visited.

- **Known Issue:**

**For Kubernetes environments, when you enable Feedback and Diagnostics for your installation, the collected data can include Kubernetes cluster names.**

To help us improve the product, you can enable Workload Optimization Manager to collect anonymized and non-confidential data as you go about using the product. However, because of the way Kubernetes discovery works for Workload Optimization Manager, the collected data includes the names of any Kubernetes clusters that you have set up as targets. We do not make use of those cluster names in any way.

If you do not want Workload Optimization Manager to collect these cluster names, then you can navigate to **Settings / Maintenance Options / Feedback and Diagnostics** and turn off the option to share anonymized usage data.

- **Known Issue:**

**Changes to a policy do not immediately show up in the user interface view of the affected scope.**

When you set the scope of the Workload Optimization Manager view to a group, you can then view the automation policies that impact the given group. If you edit a policy for that group (in Settings: Policies), and then scope the view to that group again, the policy changes do not appear in the display for that group.

The display should update within ten minutes, after the next round of incremental discovery. If the condition persists, log out of your session and log in again to update the display.

- **Known Issue:**

**For Migrate to Cloud plans, under rare circumstances the plan's actions list can show duplicate entries.**

For Migrate to Cloud plans, under rare circumstances the plan's actions list can show duplicate entries.

- **Known Issue:**

**For Azure environments, discovery does not support the Brazil Southeast region.**

For Azure environments, Workload Optimization Manager does not discover the Brazil Southeast region. Azure provides this region only to give business continuity and disaster recovery to workloads in Brazil South that require data residence.

The user interface does not display the Brazil Southeast region in any lists or charts. Also, if you do have workloads on that region, Workload Optimization Manager will not discover those workloads.

- **Known Issue:**

**For AppDynamics environments, the platform cannot discover Database Servers if the target authentication uses OAuth for credentials.**

For AppDynamics environments, Workload Optimization Manager cannot discover Database Servers if the target authentication uses OAuth for credentials.

- **Known Issue:**

**For Application Component automation policies, the user interface allows you to make conflicting settings.**

The Action Generation setting can show incorrect values that you can choose for the policy. As a result, you cannot save the policy.

- **Known Issue:**

**The user interface does not currently show the billed costs for some Azure resource groups.**

For Azure environments, when you inspect resource groups, Workload Optimization Manager does not currently show the billed costs for those resource groups.

- **Known Issue:**

**For cloud environments, under rare circumstances analysis can recommend resizing a VM to an instance type that is older and less capable than an equally priced instance type.**

Under most circumstances, when a cloud provider offers a new instance type that is meant to replace an older type, the provider offers it at a lower cost. In at least one instance we have seen a case with identical costs for the newer and older instance types. If this occurs, and capacity and cost are equal, Workload Optimization Manager cannot ensure that it chooses the newer instance type.

To work around this issue, you can create an Action Automation policy that excludes the older instance type.

- **Known Issue:**

**The All Actions chart does not include pending actions for databases or database servers.**

The All Actions chart does not include pending actions for databases or database servers.

- **Known Issue:**

**There is a memory limit for the data you can download from the All Actions chart.**

There is a memory limit for the data you can download from the All Actions chart. For example, assume you have executed many actions over time in your environment. As a result, the list of all executed actions might exceed the data limit. In that case, downloading a CSV file from the All Actions chart will fail.

- **Known Issue:**

**Under rare circumstances, the etcd.service can fail.**

Under rare circumstances the Workload Optimization Manager platform stops responding. This occurs when `etcd.service` fails. When it does occur, you should see the following error:

```
Error response from daemon: endpoint with name etcd1 already exists in network host
```

To recover from this situation, restart the docker service for the Workload Optimization Manager platform. execute the command: `sudo systemctl restart docker.service`

- **Known Issue:**

**You must use certain templates when using PLACE to set up reservations or deployments.**

When you use the **PLACE** page to set up a reservation or a deployment, you choose the templates to represent the workload you will deploy. The templates you choose must include an **Image** specification that gives the path to the VM package, and optional placement constraints.

Typically, you will use templates that are discovered through your hypervisor targets. Along with discovering resource capacities for the given VM, Workload Optimization Manager should also discover the Image specification for a given discovered template. However, in this version Workload Optimization Manager does not discover the Image descriptions. In addition, discovered templates and their image specifications are read-only. For this reason, you cannot set up placement or reservations using discovered templates.

- **Known Issue:**

**For resources that do not support Reserved Capacity, charts can show them with zero reserved capacity.**

Ring charts that show the utilization of different resources show a yellow segment whenever the Reserved Capacity for the resource is zero. For some resources there is no concept of reserved capacity, yet the ring chart still shows a yellow segment.

- **Known Issue:**

**Optimized Improvements for plans do not include hosts to provision.**

For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.

- **Known Issue:**

**In vCenter environments, you might see high storage latency, or excessive storage provision.**

In vCenter environments, you might see unusually high storage latency values or excessive recommendations to provision new storage. There is a known problem with the storage latency values that vCenter Server versions 6.5.u1x and earlier return via the API. These versions can return unusually high storage latency values.

Workload Optimization Manager considers storage latency when calculating whether to move a VM to existing storage, or whether to provision new storage. Because of this known problem, Workload Optimization Manager can incorrectly recommend provisioning storage when moves are appropriate.

If you encounter this problem, then you should create a policy that disables storage moves for VMs that are managed by vCenter Server versions 6.5.u1x and earlier. To create this policy:

- Create a VM group that contains all the affected VMs. Note that Workload Optimization Manager automatically creates a group named `VMs_vCenter` that you might be able to use.
- Create a new VM automation policy. This policy will disable storage move actions.
- Set the group that you created to be the policy scope.
- Under **Action Automation** add the `Storage Move` action and set it to `Disabled`.

- **Known Issue:**

**The Optimal Improvements chart can show incorrect data for hosts to be suspended.**

In cases where actions recommend that you suspend hosts, the Optimal Improvements chart should indicate no utilization on the hosts to be suspended. Under some circumstances, the chart can show utilization on these hosts. The result is incorrectly low values for utilization on the other hosts in the current scope.

- **Known Issue:**

For vSAN environments, under some circumstances a plan that is scoped to a datacenter can fail.

For vSAN environments, when running plans that add or replace hosts to the environment, under some circumstances the plan can show the incorrect count for hosts, and the plan can fail.

This can happen for plans that meet the following conditions:

- The plan type is Hardware Refresh, Add Workload, or Custom
- The plan scope is set to a datacenter, and it includes vSAN hosts
- The plan uses an HCI template to replace the hosts

After running, the plan shows the full count of hosts in the vSAN environment, instead of the count of hosts in the plan scope.

To avoid this situation, do not scope the plan to the datacenter.

- **Known Issue:**

**The Headroom chart for All On-prem Hosts does not always agree with the Top Clusters chart.**

The Headroom chart for All On-prem Hosts does not always agree with the Top Clusters chart.

Workload Optimization Manager generates the All On-prem Hosts headroom data in a nightly plan. When the plan runs, this data is correct. In the course of the day, this data can become stale.

To accurately track your cluster usage, you should use the Top Clusters chart.

- **Known Issue:**

**For vCenter Server environments, Workload Optimization Manager does not recognize DRS rules for VM restart dependencies that are based on ClusterDependencyRule.**

For vCenter Server environments, Workload Optimization Manager does not recognize DRS rules for VM restart dependencies that are based on `ClusterDependencyRule`.

You might be able to achieve a similar effect by expressing dependencies via `ClusterVmHostRule` or cluster affinity or antiaffinity rules.