



Workload Optimization Manager 3.4.4 Release Notes

September 23, 2022

This document describes issues that are addressed in Workload Optimization Manager 3.4.4 – Release Date: September 23, 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
3.4.4	8.6.4
3.4.3	8.6.3
3.4.2	8.6.2
3.4.1	8.6.1
3.4.0	8.6.0
3.3.7	8.5.7
3.3.6	8.5.6
3.3.5	8.5.5

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 3.4.4

Version 3.4.4

- **Support for Microsoft Azure Billing Targets**

This release introduces support for Microsoft Azure Billing targets. Workload Optimization Manager discovers your Azure billing account and related subscriptions through your Enterprise Agreement (EA) offer ID.

NOTE:

Use Microsoft Azure Billing targets for non-government accounts.

For details, see "Microsoft Azure Billing Targets" in the *Target Configuration Guide*.

- **Delete Actions for GCP Volumes**

Workload Optimization Manager can now recommend deleting unattached GCP volumes as a cost-saving measure.

In this release, Workload Optimization Manager generates delete actions for [zonal](#) (single zone) persistent disks. Use the Potential Savings or Storage Summary chart to view and execute actions.

For details, see "GCP Actions" in the *Target Configuration Guide*.

- **Red Hat OpenShift Installation Instructions**

The Installation Guide now includes instructions for installing Workload Optimization Manager on Red Hat OpenShift. Please review the installation prerequisites and steps before performing the installation.

For details, see "Installing on Red Hat OpenShift Clusters" in the *Installation Guide*.

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

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 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).

If you need to migrate from a Classic installation to one of the 3.x version families, please contact your support representative.

Configuration Requirements

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

Dynatrace Targets

Starting with Workload Optimization Manager version 3.4.2, the API token that you use when you configure a Dynatrace target must access specific scopes of the Dynatrace API V1 and V2.

If you are updating to Workload Optimization Manager version 3.4.2 or later, from a version that is earlier than 3.4.2, you must generate a new API token for the Dynatrace target configuration. Then you must enter that token in the target configuration, and validate the target.

Workload Optimization Manager uses the API token to authenticate its calls to the Dynatrace API. This token must have permission to execute GET methods via the Dynatrace API, both Version 1 and Version 2. Generate a new generic access token with these scopes:

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ▪ API V1 scopes: <ul style="list-style-type: none"> – Access problem and event feed, metrics, and topology ▪ API V2 scopes: <ul style="list-style-type: none"> – Read entities – Read metrics

NOTE:

Under some circumstances, after you update the access token, the target still fails to validate. If that occurs, take note of your configuration settings, delete the target, and configure the target again. Be sure to use the new API token that you have generated.

For more information, see the Known Issue, ["For Dynatrace targets, you must update the API Token..." \(on page 8\)](#).

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:
3.4.4	42.16
3.4.3	42.15

Product Version:	Operator Version:
3.4.2	42.14
3.4.1	42.13
3.3.7 - 3.4.0	42.12
3.3.6	42.11
3.3.4 - 3.3.5	42.10
3.3.2 - 3.3.3	42.9
3.2.6 - 3.3.1	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 latest version of the *Workload Optimization Manager 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 OperatorHub, 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.
`kubect1 apply -f operator.yaml`
 4. Validate that the Operator pod is ready.
 Execute the command: `kubect1 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 OVA and VHD Installations

For its default historical database on OVA and VHD installations, Workload Optimization Manager currently supports MariaDB version 10.5.16. This support includes comprehensive testing and quality control for Workload Optimization Manager usage of the historical database.

IMPORTANT:

Because of a known issue, *you must never use* MariaDB versions 10.5.14, 10.5.15, 10.6.7, 10.7.3, or 10.8.2.

If you are running Workload Optimization Manager installed as an OVA or VHD image, and using the database that is included in that installation, then you must use version 10.5.16. For versions of Workload Optimization Manager that you installed as an OVA or VHD *before* version 3.3.6, if you have not explicitly updated MariaDB to 10.5.16, then you must do it now.

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 is 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.

Fixed Issues

- **Fixed Issue:**

After you update to version 3.4.3, the `prometheus-server` pod remains stuck in the `ContainerCreating` state.

After you update to version 3.4.3, the `prometheus-server` pod remains stuck in the `ContainerCreating` state.

- **Fixed Issue:**

For Azure, under some circumstances the billing target does not discover RI utilization.

If the only cost in your Azure subscription is for RIs, then the Azure environment does not return an ID for the billing offer. In that case, discovery does not process the associated costs, and does not recognize utilization of the affected RIs.

- **Fixed Issue:**

For environments with many ephemeral entities, historical processing of topology data can suffer performance issues.

Workload Optimization Manager saves historical data for entities in your topology. However, for environments with many ephemeral entities (entities with limited lifespans), the topology processor can experience performance issues. In this case the user interface can display incomplete data.

If you experience this issue, you should see errors that contain:

```
ERROR [realtime-pipeline-runner] [TimeSlotEditor] : Maintenance failed for...
```

- **Fixed Issue:**

Under rare circumstances, the user interface can display Reassemble Event notifications.

Under rare circumstances, Workload Optimization Manager can display Reassemble Event notifications in the user interface. Logging can also show messages that contain:

```
ERROR [PercentileEditor] : Reassemble failed for...
```

This can happen when the process gets a null value for a connection, but tries to close that connection.

- **Fixed Issue:**

For MySQL targets, under rare circumstances discovery can fail.

For MySQL targets, under rare circumstances discovery can fail. This can occur when discovery attempts to access empty data sets.

- **Fixed Issue:**

Dynamic groups of unattached volumes can fail to apply action policies to their members.

Assume you use the **Days Unattached** filter to create a dynamic group of unattached volumes. If you assign this group a policy to delete unattached volumes, the policy does not take effect.

- **Fixed Issue:**

For VMware Horizon environments, under some circumstances discovery can fail to complete.

This can happen when your environment contains nested groups, or groups that share members.

- **Fixed Issue:**

For vCenter Server, under some circumstances discovery can fail to build the supply chain.

In a vCenter Server environment, your entities can use tagging in the VCTAGS namespace that gives multiple keys with the same property key. In this case, discovery fails with an error that contains:

```
StitchingManager : Unable to apply stitching operation
StringsToStringsDataDrivenStitchingOperation...
```

- **Fixed Issue:**

For Container Spec action policies, you cannot use a policy that passes the action to orchestration.

For Container Spec policies, you can define a policy that sends the action to ServiceNow for approval. However, Workload Optimization Manager does not apply the policy to your Container Spec scope.

- **Fixed Issue:**

The Data Extractor component can flood the log with messages.

Under some circumstances, if the extractor encounters an error common to all entities in your environment, it can flood the log with unnecessary repeated messages. The log now lists the error once, and provides a count of the incidents.

- **Fixed Issue:**

For Google Kubernetes Engine (GKE) environments, Workload Migration plans can fail to place pods.

When running a Workload Migration plan that places pods on a GKE cluster, the plan can fail to place the pods.

- **Fixed Issue:**

For HPE OneView, under certain circumstances, target validation can fail.

For HPE OneView, under certain circumstances, target validation can fail with an error that is similar to:

```
MalformedJsonException: Invalid escape sequence at line 1...
```

Known Issues

- **Known Issue:**

Updating from earlier versions can impact your vCPU Scaling policies.

When you upgrade from the 3.3.x version family to the 3.4.x version family, the update can impact the vCPU Scaling for on-prem VMs that is defined in the **Virtual Machine Defaults** policy.

If, in the earlier version, you changed the **Increment Constant for vCPU** setting for the **Virtual Machine Defaults** policy, you should *review the new documentation for setting vCPU Scaling policies before you perform a product update.*

After you perform a product update, you should consider the following:

- Use a new vCPU Scaling Control to achieve the desired vCPU resizing outcome
- Use the **vCPU Resize Min Threshold (in cores)** setting to prevent VMs from resizing below a desired number of vCPUs
- Use the vCPU Scaling Control **MHz legacy behavior** to continue using the custom default setting in MHz

• **Known Issue:**

Under some circumstances, the Storage Summary chart can show an incorrect number of volumes.

If you compare the total count of volumes in the Storage Summary chart with the total count of volumes in Tiers, the counts can be inconsistent.

• **Known Issue:**

For Dynatrace targets, you must update the API Token that you use for Target Configuration so that it supports V1 and V2 of the Dynatrace API.

The Target Configuration Guide does not specifically state that you must update the API token for existing Dynatrace targets, and it does not list the new permissions the token needs for each API version.

To configure a new Dynatrace target, generate an access token with the permissions listed below.

If you are updating to Workload Optimization Manager version 3.4.2 or later, from a version that is earlier than 3.4.2, you must generate a new API token for each existing Dynatrace target. Then you must enter that token in the target configuration, and validate the target.

NOTE:

Under some circumstances, after you add a new token to the target, it still does not validate. If that occurs, take note of your configuration settings, delete the target, and configure the target again. Be sure to use the new API token that you have generated.

The Dynatrace API token must include permissions to access V1 and V2 of the Dynatrace API. Generate a generic token with these scopes:

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> – API V1 scopes: <ul style="list-style-type: none"> • Access problem and event feed, metrics, and topology – API V2 scopes: <ul style="list-style-type: none"> • Read entities • Read metrics

• **Known Issue:**

Cloud VMs might scale to very large instance types even if there are smaller instance types available for scaling.

When a cloud VM with a specific disk count and disk type applies a policy that enables 'instance store aware scaling', Workload Optimization Manager might recommend scaling the VM to a very large instance type, even if there are smaller, less expensive instance types that can adequately meet the VM's resource requirements.

To avoid this issue, disable instance store aware scaling and restrict the VM to its current instance family. For example, if an AWS VM is currently running the `i3.2xlarge` instance type, specify the `i3` instance family as a scaling constraint for the VM.

- **Known Issue:**

If you use Embedded Reports, your dashboards might be empty after updating to this version.

To support Embedded Reports, Workload Optimization Manager loads historical data into a TimescaleDB service, and displays that data in dashboards and charts that are powered by the Grafana observability platform. Prior to version 3.4.1, Workload Optimization Manager used the Enterprise version of Grafana, and included an Enterprise license for customers who paid for the license to use Embedded Reports.

Starting with version 3.4.1, Workload Optimization Manager no longer has an agreement to use the Enterprise version, and ships the open-source Community version instead. After you update to version 3.4.1, Grafana automatically switches to the Community version and you should be able to use Embedded Reports as usual.

However, under rare circumstances, you might notice that your dashboards are empty and a banner with this error message displays:

```
db query error: pq: password authentication failed for user "query"
```

To resolve this issue, follow these steps:

1. Delete your Embedded Reports data source.
 - In the Grafana user interface, navigate to **Configuration / Data Sources**.
 - Click to select the Workload Optimization Manager timescale data source.
 - In the screen that appears, click **Delete** to delete the Workload Optimization Manager timescale data source.
2. Restart the extractor component.

When you restart the extractor component, it loads the Grafana Community version and reconfigures the data source.

- Get the full pod name for the extractor pod by executing the command:

```
kubectl get pods -n {YourNamespace} | grep ^extractor
```

For example, execute:

```
kubectl get pods -n turbonomic | grep ^extractor
```

The full pod name should be similar to:

```
extractor-5b86976bc8-vxwz4
```

- Delete the extractor pod.
Deleting the pod triggers it to restart. Using the example name above, you would execute the command:

```
kubectl delete pod -n {YourNamespace} {YourExtractorPodName}
```

For example:

```
kubectl delete pod -n turbonomic extractor-5b86976bc8-vxwz4
```

For further assistance, contact your support representative.

- **Known Issue:**

The user interface shows incomplete Azure billed costs and metrics because of an issue with the Azure API.

When attempting to discover Azure billing targets configured with partitioned cost exports on the first day of any month, Workload Optimization Manager is unable to find the directory containing the export files within a customer's Azure Storage account. The export files exist, but the Azure API does not return the correct storage location of that day's cost export files. As a result, the Workload Optimization Manager user interface may not reflect any billed cost that Azure exported on the first of the month, and the metrics dependent on that billed cost. Note that the billed cost is not limited to expenses incurred on the first day of the month.

In the logs, the following one-time message with an INFO severity appears after the first discovery attempt.

```
Download storage blob failed: no blob found from prefix: ${prefix}
```

- **Known Issue:**

In a Nutanix environment, a Replace Hosts plan can fail to place VMs.

You can configure a Replace Hosts plan on a Nutanix cluster to replace the hosts with HCI templates. However, the plan will fail to create the HCI hosts, and will result in unplaced VMs.

- **Known Issue:**

When you view the details table for the Top Accounts chart, the Actions Taken and Saved By Actions columns can fail to show current data.

When viewing the Top Accounts chart, you can click **Show All** to see a details table. Under some circumstances, this table can fail to update with new data in the **Actions Taken** or **Saved By Actions** columns. For example, these columns can show zero even if you have taken actions that reduced costs.

To see current actions data, display the Executed Actions chart.

- **Known Issue:**

Hardware Refresh to replace hosts with HCI templates can fail to place workloads.

When running a Hardware Replace plan, the plan can fail to place workloads onto HCI hosts. The plan correctly places workloads if the plan scope is in a hyperconverged environment. If the scope is *not* in a hyperconverged environment, then you must scope the plan to an entire cluster, and you must configure the plan to replace all the hosts in the cluster with HCI templates.

- **Known Issue:**

Under specific conditions, an update of the platform can cause it to lose all of its target configurations.

When updating Workload Optimization Manager from a version earlier than 3.3.3, if the username for any user accounts includes a % character, the update will fail to include your target configurations.

If this occurs when you update, contact your support representative for assistance.

NOTE:

You should always back up your installation before you run an update.

- **Known Issue:**

For AWS environments that have RI discount sharing turned off, analysis does not manage RI coverage and utilization.

In AWS, you can turn off RI discount sharing for specific accounts. These accounts will not share any discounts with other accounts. Workload Optimization Manager does not recognize RI coverage or utilization for these accounts. For example, the RI Coverage and RI Utilization charts will show zero values.

If you encounter this situation, contact your support representative for a possible workaround.

- **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 Kubernetes 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 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.