



Turbonomic 6.4.3 Release Notes

October 4, 2019

This document describes issues that are addressed in Turbonomic 6.4.3 – Release Date: October 4, 2019. All builds are cumulative. Applying 6.4.3 onto any release of Turbonomic v5.9 or later will include all previous fixes. Please see the Turbonomic documentation for earlier versions of the Release Notes:

<https://greencircle.vmturbo.com/community/products/pages/documentation>

For any questions, please contact Turbonomic Technical Support at support@turbonomic.com, or open a ticket at:

<https://greencircle.vmturbo.com/support>

For Updates to the 6.4 Family of Turbonomic Releases

To ensure performance of your installation, we have increased the minimum requirement for memory on the VM that runs Turbonomic. The new minimum requirement for memory is 32 GB. For more information about minimum requirements, please see the *Turbonomic Installation Guide*.

What's New for Version 6.4

The 6.4 family of Turbonomic releases includes the following new features:

- Planning Improvements

- Overhaul of the Planning Workflows and Functionality

This release includes a significant reorganization of the Plan user interface. It should be easier to use wizards to configure plans, and easier to configure custom plans. In addition, you should find it easier to review and interpret the plan results. This overhaul includes improvements to the Cloud Comparison chart that make it easier to see how Turbonomic calculates savings and investments on the cloud.

See "Setting Up Plan Scenarios" in the *User Guide*.

- Plan Migrations of On-Prem Databases to Azure SQL Managed Instances

When migrating your on-prem environment to the Azure cloud, Turbonomic identifies SQL Server instances. You can configure the plan to migrate those database instances to Azure SQL Managed Instances.

See "Migrate SQL Database to Azure Plan" in the *User Guide*.

- Integration of Plan Results with the Azure Migrate Portal

After you plan a migration from on-prem to your Azure environment, you can upload the plan results to the Azure Migrate portal.

See "Uploading Plan Results to Azure Migrate" in the *User Guide*.

- Improved Support for Azure Environments

In addition to enhanced plans for migrations to Azure, this release adds support for:

- Azure Enterprise Agreements

Turbonomic now recognizes Microsoft Azure Enterprise Agreements (EAs). When you configure an EA target and set the EA key to your Azure targets, Turbonomic uses that richer pricing information to calculate workload placement. For more information, see "Azure Enterprise Agreements" in the *User Guide*.

This release makes it easier to add Azure Accounts as targets. Starting with this release, you can add a single Service Principal target, and Turbonomic discovers the individual accounts related to that Service Principal. For more information, see "Azure Enterprise Agreements" in the *Target Configuration Guide*.

- Azure Reserved VM Instances

If you add an Azure EA target, the pricing includes costs for RI VMs. Turbonomic uses this information to recommend RI Buy actions, and to calculate the savings you should see as a result.

- Cloud Native and Container Improvements

This release includes

- Consistent Scaling Actions for Containers

Turbonomic now discovers and groups the containers for a given service. It then enables a policy that maintains consistent resizing for the group of containers. As a result, Turbonomic generates actions to resize all the members of that group to the same size, such that they all support the top utilization of each resource commodity in the group. Executing the actions will then update the container deployment.

See "Analysis Policies: Containers" in the *User Guide*.

- Horizontal Cluster Actions

For Kubernetes environments where the Cluster API is present, this release introduces support for Kubernetes Node Provision actions. Turbonomic discovers whether the Cluster API is present. To set up elastic Kubernetes clusters, you can set the provision actions to Manual or Automated.

- Capacity Planning for Cloud Native

You can configure custom plans to explore what happens if you:

- Scale demand in or out horizontally (add or remove containers or container pods)
- Scale demand vertically (increase or decrease utilization for groups of containers or pods)
- Change cluster resources (for example, take down nodes or change templates)
- Ignore constraints in your container environment
- Change action policies and apply specific placement policies to the scope of your container environment

For more information, see "Container Utilization Plan" in the *User Guide*.

- Recognition of Requests and Reservation of Resources

In Kubernetes environments, scheduling takes into account guaranteed resource requests. You can now see how much request capacity is currently utilized for a given node (VM), Pod, or namespace, alongside a view of how the limits capacity is utilized. Turbonomic uses this information when calculating actions to manage pod placement, and to manage underlying resources. See "Supply Chain – ContainerPod" in the *User Guide*.

- Native Integration with Service Now

This release introduces *Turbonomic Actions*, a ServiceNow application that you install via an update set. Turbonomic Actions integrates Turbonomic with your ServiceNow account to log the actions Turbonomic generates, and to defer these actions to your ServiceNow workflow for approval. Upon approval, Turbonomic executes the actions during the scheduled maintenance window, and logs the result in ServiceNow.

For information about setting up policies for these actions, see "Action Orchestration" in the *User Guide*. To get the Turbonomic Actions application, contact your sales representative.

- Improved Action Scheduling for Automation Policies

Turbonomic includes improvements to action scheduling. It is now easier to set up action schedules. Also, if Turbonomic recommends an action outside of the schedule window, it queues the action for later execution at the scheduled time. If the action is still valid at that time, Turbonomic will execute it. See "Working With Schedules" in the *User Guide*.

- Integration with VMware Horizon VDI

To support virtual desktop infrastructures, Turbonomic has added new entities to the Supply Chain:

- Business User – A user account that is entitled to launch one or more active VDI sessions
- Desktop Pool – A collection of desktops that support users
- View Pod – A logical grouping of desktop pools, that provides a capacity for active sessions

For more information, see these entity types in "Supply Chain of Entities" in the *User Guide* and see their policy settings in "Analysis Settings" in the *User Guide*.

- Improved vSAN Support for Hyperconverged Environments

For environments that use hyperconverged infrastructure to provide storage on vSAN, Turbonomic can represent the storage that is provided by a cluster of HCI hosts as a single storage entity. Analysis can generate actions to scale Storage Amount, Storage Provisioned, and Storage Access in the vSAN. See "vSAN Storage" in the *User Guide*.

You can use HCI Host templates to run plans that add capacity to your vSAN. For more information, see "HCI Host Template Settings" in the *User Guide*.

- Custom Reports

To support custom reports, you can now create a Custom Report Template that executes your own SQL queries into the Turbonomic database. Once you create the custom template, you can generate and view them, schedule them to generate a copy at the times you specify, and set up subscriptions to them, the same as you can with the standard reports. For more information, see "Reports: Viewing Historical Data" in the *User Guide*.

- OpenStack Queens Support

For use in OpenStack environments, Turbonomic is now certified to support OpenStack Queens.

Update Recommendations

You can apply this update to any GA version of Turbonomic from version 5.9 or later.

NOTE:

If your Turbonomic installation is running on an older version, or if it is running on openSUSE, contact Turbonomic Technical Support to confirm your update path.

If your server is able to connect to the Internet, you can apply the update through the online process by going to **Admin > Maintenance > Software updates > Update**.

Offline Update Links

Turbonomic 6.4.3 is available as an offline update. Please see the Green Circle article:

[How To Perform an Operations Manager "Offline Update" - Latest Links Included](#)

NOTE:

For this version of Turbonomic we recommend an allocated disk storage capacity of 500GB.

Related Green Circle Articles

- Offline and Online Update Instructions:

<https://greencircle.vmturbo.com/docs/DOC-1649>

- Release Notes and Product Documentation:

<https://greencircle.vmturbo.com/community/products/pages/documentation>

Configuration Requirements

For this release of Turbonomic, you should satisfy the following configuration requirements.

Updating the Tomcat Server

There are circumstances when you might choose to upgrade the Tomcat server on Turbonomic to a later version. In this case you must copy a local configuration file to the tomcat installation.

After you update the Tomcat server:

- Copy the file `/usr/libexec/tomcat/server.local` to `/usr/libexec/tomcat/server`
- To ensure that this server configuration file is executable, perform the command: `chmod 755 /usr/libexec/tomcat/server`

Transport Layer Security Requirements

Starting with version 5.4, by default Turbonomic requires Transport Layer Security (TLS) version 1.2 to establish secure communications with targets. Most targets should have TLSv1.2 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 Turbonomic 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 TLSv1. If your NetApp target suddenly fails to validate after installing Turbonomic 5.4 or later, this is probably 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 Turbonomic Technical Support.

Enabling HTTP and HTTPS Proxies

Turbonomic supports the use of HTTP and HTTPS proxies for internet communication. However, you must edit the Tomcat Server configuration file to add the required system variables.

The file you must edit is on your Turbonomic server at `/usr/libexec/tomcat/server`

In this file, search for the `OPTIONS` statement. It should appear in the config file similar to the following:

```
FLAGS="$( $JAVA_OPTS) $CATALINA_OPTS "
```

```
OPTIONS="-Dcatalina.base=$CATALINA_BASE ..."
```

Add the following flags to the `OPTIONS` statement, giving values for your proxies:

```
-Dhttp.proxyHost
```

```
-Dhttp.proxyPort
```

```
-Dhttps.proxyHost
```

```
-Dhttps.proxyPort
```

```
-Dhttp.proxyUser
```

```
-Dhttp.proxyPassword
```

```
-Dhttps.proxyUser
```

```
-Dhttps.propyPassword
```

The resulting `OPTIONS` statement should be similar to the following:

```
OPTIONS="-Dcatalina.base=$CATALINA_BASE \  
-Dcatalina.home=$CATALINA_HOME \  
-Dhttp.proxyHost=111.10.10.123 -Dhttp.proxyPort=123 \  
-Dhttps.proxyHost=112.10.10.123 -Dhttps.proxyPort=456 \  
-Dhttp.proxyUser=user -Dhttp.proxyPassword=password \  
-Dhttps.proxyUser=user -Dhttps.propyPassword=password"
```

Note that the values you provide for this file must match the values you provide when specifying a proxy in the Turbonomic user interface. After you make these changes, restart the Tomcat server.

For further assistance, contact Technical Support.

Enabling HTTP to HTTPS Redirects for the REST API

Starting with version 5.9.1, Turbonomic redirects http requests to the user interface and the REST API over to HTTPS. For the user interface to display, this redirect must be in place.

If you are updating from a version of Turbonomic that is earlier than 5.9.1, then you must restart the httpd service on the Turbonomic server. A restart enables this redirect after such an update.

To restart the httpd service, use the following command: `service httpd restart`

For more information, contact Technical Support.

Enabling IOPS and Network Monitoring for OpenStack Mitaka

The Target Configuration Guide gives instructions to connect to OpenStack targets. However, if you are running OpenStack Mitaka, you must perform additional configuration on the Mitaka platform to enable IOPS and Network data collection from Physical Machines. For those instructions, please see the Green Circle article, [Enabling OpenStack PM Metric Collection](#).

SMI-S Provider Versions for EMC VNX and EMC VMAX Storage Solutions

To connect to EMC VNX and VMAX disk arrays, Turbonomic uses EMC SMI-S providers that have the given disk arrays added to them. Note that VNX and VMAX support different versions of SMI-S Providers:

- **VNX**
For VNX and VNX2 arrays, use SMI-S version 4.6.2, based on Solutions Enabler 7.6.2. We have verified Turbonomic control of VNX block storage using SMI-S version 4.6.2 as a target.
- **VMAX**
For VMAX arrays, use SMI-S version 8.1, which is included in Solutions Enabler 8.1 – We have verified Turbonomic control of VMAX storage arrays using SMI-S version 8.1 as a target.

Fixed Issues

- **Customer Issue 107860**
Fixed: In Hyper-V environments that include multiplexor devices, Turbonomic can overestimate the network throughput capacity on Hyper-V host machines.
- **Customer Issue 107751**
Fixed: Under some circumstances, Turbonomic can incorrectly identify storage as wasted, and recommend deleting that storage even though there is an active VM on it. This can happen when the discovery process requests VM data from the Azure target, and the target returns an error. In that case, Turbonomic does not discover the affected VM, and does not recognize that the VM is active on its storage.

- **Customer Issue 107754**

Fixed: For Azure environments, under rare circumstances after updating from the 6.3 version family to the 6.4 version family, discovery can fail for given Azure targets.

- **Customer Issue 107553**

Fixed: For large Microsoft Virtual Machine Manager environments, under some circumstances discovery takes an excess amount of time.

- **Customer Issue 107749**

Fixed: Assume you successfully creates a reservation (in the Place page), if you then deleted the reservation template Turbonomic cannot retrieve the reservation or display it in the user interface. This is normal behavior. However, when this occurs, Turbonomic throws a Null Pointer error. The software should handle this situation more gracefully.

- **Customer Issue 107701**

Fixed: When running a Migrate to Cloud plan, if the plan handles VMs that refer to the Azure `Standard_E20_v3` template, the plan shows incorrect results.

- **Customer Issue 107741**

Fixed: When you set the scope of an Optimize Cloud Plan to be a group of VMs, the Optimize Cloud Plan wizard provides two Optimization Settings that are not valid for a scope of VMs:

- Fully optimize cloud services with existing and new reserved instances
- Optimize cloud spend with the purpose of new reserved instances

For these options to be valid, the plan scope must be set to AWS accounts or Azure subscriptions. To fix this issue, the wizard must not provide these options when the scope is set to a group of VMs.

- **Customer Issue 107612,107679**

Fixed: In public cloud environments, when you set the scope to a single cloud subscription or account, the pending actions list loads all the actions for that scope, and does not page the actions request. If the list of actions for that scope is large, this can result in unacceptable performance when displaying the user interface.

- **Customer Issue 107663**

Fixed: In large public cloud environments, under some circumstances it is possible that Turbonomic will not recommend RI Buy actions even though they are appropriate.

- **Customer Issue 107523**

Fixed: For large environments that you have configured to use Merge placement policies, under some circumstances Turbonomic does not recommend Shared Nothing moves as it should. This can happen when the performance for incremental discovery is poor.

- **Customer Issue 107619**

Fixed: For environments that have only public cloud targets, health checks post warnings for on-prem hardware statistics (`pm_stats_by_hour` and `ds_stats_by_hour`). In a cloud-only environment the health check should not post these notifications.

- **Customer Issue 107543**

Fixed: The click behavior for the Top Accounts: Show All chart is not correct. To expand or collapse entries in this chart, if you do not click exactly on the chevron icon, the list item does not expand or collapse. Instead, it navigates your view to a different scope. This chart should be easier to use with mouse gestures.

- **Customer Issue 107468**

Fixed: For vCenter Server environments that use hyperthreading, it is possible for a hyperthreaded VM to have a greater vCPU capacity than the CPU capacity that Turbonomic calculates for the given host. In that case, the user interface can show inconsistent values for vCPU on the VM.

- **Customer Issue 105911**

Fixed: In the Place View, when you execute a placement (click **Find Placement**) but the placement fails, the user interface displays a notification that tells you the placement failed, but it does not give the reason for failure. Instead, it tells you to look at the log. The notification should include the reason for failure.

- **Customer Issue 107417**

Fixed: After moving a host out of maintenance mode, under some circumstances it takes too long for the host to re-populate with workloads. This can happen if the utilization history of the host still retains peak values that correspond to the workloads the host had before it went into maintenance.

If not enough time has passed between when the host was put into maintenance mode and when the host is taken out of maintenance mode, then we will slowly update the `histUtilizations` based on the used weight (50%) and peak weight (99%). At 0.99 the peak weight is extremely resistant to updates. This fix clears the historical utilization and uses the current utilization values when the host comes out of maintenance mode. It also lowers peak weight from 99% down to 90% to update historical peak values more quickly.

- **Customer Issue 107416**

Fixed: In AWS environments that include billing families, if Turbonomic encounters an error when discovering the AWS environment, it is possible that Turbonomic will recommend unnecessary reconfigure actions for workloads in the billing family to address `VMPMAccessCommodity` issues.

- **Customer Issue 107399**

Fixed: The display is confusing when you compare values in the Capacity and Usage chart to the values in the Multiple Resources chart. For the same metric (say Storage Latency), the Capacity and Usage chart shows amount utilized for the metric first, while the Multiple Resources chart shows the percentage utilized. If you compare the two, they seem to show different values for the same metric.

- **Fixed:** For WebLogic environments, Turbonomic can recommend actions to resize thread pools. However, in WebLogic the thread pools resize dynamically, and users or clients to WebLogic cannot execute resize actions. Turbonomic should not recommend thread pool resizes for WebLogic.

- **Customer Issue 107364**

Fixed: In 3PAR environments where storage capacity is a Petabyte or greater, target validation fails with an error that expects an integer value.

- **Fixed:** When viewing the Optimize Cloud plan results, the Unattached Volume Tier Breakdown chart is missing account information.

- **Customer Issue 107228**

Fixed: Setting a scheduled recurrence when the user interface and the Turbonomic server are on different time zones can have unintended effects.

For example, assume the user interface is on Mountain Standard time (MT) and the server is on Eastern Daylight time (ET). Further assume that on a Friday you schedule a recurrence for Fridays at 10:00 pm MT. 10:00 pm Friday for the user interface is 12:00 pm for the server. Because of the date change at 12:00 pm, the server skips the first scheduled occurrence.

- **Customer Issue 107607**

Fixed: In AWS environments, Turbonomic can recommend moving root EBS volumes to invalid storage tiers.

- **Customer Issue 106816,106827,106839,107146,107609**

Fixed: Under some circumstances after restarting Turbonomic, when you navigate away from the Home Page and navigate back to it again, actions appear to be removed from Pending Actions list. After another cycle of incremental discovery the actions reappear.

- **Customer Issue 106661**

Fixed: In public cloud environments, under some circumstances the details of an action to resize a VM can appear to change unexpectedly.

Assume a VM on Template_A. If Turbonomic recommends resizing the VM to Template_B, but later recommends resizing it to Template_C, then the action entry will appear to change. To fix this issue, Turbonomic must generate a unique action entry for the second resize.

Known Issues

- The Monthly Savings or Investments Totals and the Cumulative Savings charts do not include RI Buy actions in their calculations.

- **Customer Issue 107562**

Under some conditions, Turbonomic changes the action mode of an action from MANUAL to RECOMMENDED. However, this effect is not described in the documentation.

Turbonomic makes this change as a safeguard against executing actions that the underlying infrastructure cannot support. For example, assume you have VM move actions set to MANUAL. Then assume Turbonomic analysis wants to move a VM onto a host that is already utilized fully. In this case, there would be other actions to move workloads *off* of the given host to make room for this new VM. However, because moves are MANUAL, the host might not be properly cleared off yet. In that case, Turbonomic changes actions to move workloads *to* the host from MANUAL to RECOMMENDED.

- Under some circumstances, the user interface temporarily clears all actions out of the pending actions lists. After the next round of incremental discovery, the actions will reappear in these lists.

For example, this can occur when Turbonomic discovers a new Azure Scale Set. When it discovers the new Scale Set, Turbonomic creates a new group and applies a new policy to that group. As it does this, it must clear all actions and then recalculate them to account for the new group and policy. After the next round of incremental discovery, Turbonomic will populate all the pending action lists again.

- When editing automation policies, for Default Policies you can set the Rate of Resize for entity scaling. However, you cannot make a Rate of Resize setting for a scoped policy. The user interface includes a control for Rate of Resize in scoped policies, but any settings you make for a scoped policy will not take effect.
- When you run an Optimize Cloud plan, for some scopes of Azure or AWS zones and regions the Unattached Volume Tier Breakdown chart shows `No Data Available`.
- Cost data for Storage is not included in the Cloud Cost Comparison chart in the Cloud View on the Home Page for a scoped user account, even if the user's scope includes Azure resource groups. If this occurs, that same user can use the Supply Chain to navigate to a scope of hosts, and that scope will show storage costs.
- In vCenter Server environments, charts can show that a Virtual Datacenter (VDC) uses resources at more than 100% of capacity.

The utilization metrics that vCenter returns to Turbonomic for a VDC include utilization of resources that are reserved for vCenter overhead. However, the capacity metrics that Turbonomic discovers do not account for these reserved resources. As a result, it is possible Turbonomic shows that the VDC consumes more than 100% of capacity.

- The *Unattached Volume Tier Breakdown* chart shows how many unattached storage tiers or volumes exist in your storage and storage groups. This chart is designed to display in the Optimize Cloud plan results.

The user interface displays this chart for real-time views, and it enables you to add this chart to dashboards. However, this chart only shows real-time data for storage devices and storage groups. You should only rely on this chart when viewing Optimize Cloud plan results, or when your environment is specifically scoped to one or more storage devices.

- For running plans in Kubernetes or Cloud Foundry environments, you should never turn on the **Ignore Constraints** option. If you do, the plan can generate invalid results.
- In Azure environments, a subscription can use locked storage or locked resource groups. For such subscriptions, Turbonomic discovers incomplete data. Locked resources affect Turbonomic discovery in either of these scenarios:
 - A locked resource group

Turbonomic discovers all the entities in the resource group, but does not discover the resource group itself. For example, in the Top Accounts chart, the Resource Groups field will show no resource groups for a subscription that has a locked resource group.
 - Locked storage

Turbonomic discovers all the entities in the resource group except the locked storage. It also discovers the resource group.
- In Azure environments, when you run a Migrate to Cloud plan, the plan results do not include complete information for license costs for the migrated workloads.

For cost calculation, the migration plan does not include the license cost for the migrated workloads. In addition, the plan results show that the migrated workloads run on the OS platform that they ran on when on-prem. For Linux platforms, the plan results indicate Linux and do not indicate the SE Linux that Azure provides.

- For Azure environments, Migrate to Cloud plan results can show Unplaced Workloads without giving a clear explanation why.

In Azure environments, you can create a Migrate to Cloud plan that migrates Microsoft SQL Server databases to Azure as Managed Instances (a PaaS migration). However, if you select an Azure region that does not support the Managed Instance service, or does not support the instance types you have selected, then the plan cannot calculate a placement for the affected instances. In this case, the plan simply says it has unplaced workloads.

- As you run plans, you can generate duplicate plan names in the list on the Plan Management page. For duplicate names, Turbonomic should append the name with an integer to show the count of duplicate names.

This can happen if you give a plan a duplicate name before you run it. For example, Turbonomic provides a default name for each plan type. If you run multiple instances of that plan type, and do not change from the duplicate plan name, this problem will arise. To work around this problem, be sure to provide your own name for your plans.

- In AWS environments, under certain circumstances VM resizing can fail. If the restart of the VM initially fails, Turbonomic waits 30 seconds and tries to restart again. Turbonomic will try to restart up to four times. If the restart still fails, Turbonomic assumes the VM cannot start up with the new template, and it restarts the VM with the old template.

- **Customer Issue 105693**

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

To generate data for the headroom chart for All On-prem Hosts, Turbonomic uses data that is generated for the deprecated Classic user interface. The purpose of this data does not match the purpose of the chart. As a result the values in the chart do not always agree with the values in other related charts. Turbonomic 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.

- When you specify a schedule, you must be aware of issues for Daylight Savings Time changes. For example, assume you create two schedules, one that starts at 7:00 am and lasts 12 hours, and another schedule that starts at 7:00 pm and lasts 12 hours. If you apply those schedules to policies on the same scope, then the policies can experience a scheduling conflict as Daylight Savings time changes.
- For vCenter Server environments, do not create your own automation policies for clusters that are affected by vCenter HA settings.

In vCenter Server environments, Turbonomic discovers HA cluster settings, and imports them as automation policies for the affected clusters of entities. If you create a different automation policy for the same cluster, then you can get unexpected results.

For example, assume the cluster MyHosts has been configured to support HA in the event of host failure. Then, assume you create an automation policy for the MyHosts cluster that sets maximum utilization of vCPU to 75%. In that case, you can see unexpected results for Turbonomic analysis, and for the display of the affected settings.

To avoid this problem, you should understand the vCenter settings for the clusters in your environment. For any clusters that include HA settings in vCenter, you should ensure there are no user-created automation policies for the same clusters.

- For public cloud, if you have private cloud targets in your inventory (for example, VMM or vCloud Director) then Turbonomic can display private cloud entities in lists of AWS billing families or Azure subscriptions. For example, when configuring a Migrate to Cloud plan, you can configure pricing for the plan. For that step, the user interface lists the AWS billing families and Azure subscriptions that are available for the plan scope. However, this list also includes entities from the private cloud environment.

IMPORTANT:

Do not select these private cloud entities when they appear in lists for public cloud settings. If you execute public cloud functions with these private cloud entities, you might get unexpected results.

- When upgrading from a version earlier than 6.3.6, the upgrade can lose some policies that include user-created groups in their scopes. This is because the way Turbonomic manages group identities has changed. If you lose such policies, then you should create new groups and create new versions of those policies.
- When using the Classic user interface, under some circumstances the Inventory View does not update, or display content. If this occurs, perform a manual restart of the Turbonomic platform (restart the tomcat server), and the Classic user interface should perform normally.
- In AWS or Azure environments that include multiple billing families, if you scope the Turbonomic view to more than one billing family, the user interface does not display the billing family information correctly. For the **List of Billing Families** tab, it shows a count of all the accounts in the associated billing families. Also if you display that tab, the list can be empty, showing no billing families to inspect.

To work around this issue, create a group that includes all the accounts for the given billing families. Then you can scope to that group and inspect to member entities.

- For targets that attach to vCenter Server 6.7 and later, datastore browsing for VSAN can fail to ignore files that you specify in your directory ignore patterns. If you use the default ignore patterns, datastore browsing will ignore the specified files. But if you specify custom ignore strings, for vCenter 6.7 and later Turbonomic can report files you chose to ignore as wasted storage.

- Customer Issue 105298**

In AWS environments, to scale an instance Turbonomic can recommend changing a current instance to a different instance type. Note that AWS maintains a limit of EC2 instances per region.

Turbonomic does not currently get reliable information about the limits for specific instances, compared to the number of instances currently online. As a result, Turbonomic can recommend actions to move a workload to an instance that exceeds the limit. In this case, the resize action will fail.

- Customer Issue 104680,107464**

In Hyper-V environments, under some circumstances a host in a cluster can have multiple D: drives. In this case the target (Hyper-V or VMM) does not return valid data for that host, and discovery for the cluster fails. Because this is an issue of incomplete or invalid data returned by the target hypervisor, there is no known workaround that you can perform through Turbonomic.

If you experience incomplete host cluster discovery in Hyper-V, and you suspect this issue, please contact your Hyper-V support for possible ways to reconfigure your environment.

- For running plans in Kubernetes or Cloud Foundry environments, the option to disable Provisioning and Suspension of VMs has no effect. Container plans always enable Provisioning and Suspension of VMs that host containers or container nodes.

- **Customer Issue 104812,104952**

In Hyper-V environments, you must be sure that all the VMs have unique IDs.

Hyper-V supports the export of a VM, so that you can create exact copies of it by importing those exported files. The feature includes the option to choose the Import Type. The **Copy** import type creates a new unique ID for the imported VM. When importing VMs in your environment, you should always use the **Copy** import type.

Turbonomic uses the unique ID to discover and track a VM. If your environment includes multiple VMs with the same ID, then discovery will assume they are the same VM. As a result, the counts for VMs will be incorrect.

- Storage costs, as reported in the Cloud Cost Comparison chart, include costs incurred from Azure boot diagnostics storage at the normal rate for unmanaged disks.

The Cloud Services Storage cost for Azure targets, which is derived from the Azure Billing API, does not currently include costs from boot diagnostics storage. Documentation for Azure boot diagnostics storage does not specify whether it is charged at the normal unmanaged disk pricing level or at a reduced rate. We are currently seeking clarification from Microsoft.

- **Customer Issue 104265**

Starting with version 6.3.4, Turbonomic corrects a Swagger validation issue. The corrected Swagger now supports integration with Swagger development frameworks such as IBM API Connect.

To access the corrected swagger data, if you have updated from an earlier version of Turbonomic then you must restart your apache HTTPD server. Then you can access it via the approved URL. For example, to access the Turbonomic Swagger-UI, navigate to `https://[Your_Turbonomic_IP]/apidoc/v2`.

You can still access the swagger via the deprecated path (does not include an API version in the path). For example, you can navigate to `https://[Your_Turbonomic_IP]/apidoc` to see the Swagger-UI. This will work even if you do not restart the apache HTTPD.

For API developers, note that you can use the API to install multiple license files in a Turbonomic instance. However, you cannot use Swagger to install multiple license files. This limitation exists for users of Swagger-UI, and for developers using the Turbonomic swagger through a development framework.

- For targets that you have configured and validated, it is possible that changes to the configuration of the remote service can cause the Turbonomic connection to it to fail. However, the user interface still shows the connection as valid.

For example, you can configure a target to discover a scope of databases via a given port. Then if the access port for one of the databases changes, Turbonomic will no longer be able to connect to that database. As a result, subsequent discovery and monitoring of that database will fail. However, the user interface will still show it as a valid target.

- In Migrate to Cloud plans, Turbonomic identifies workloads that can run on RIs, and recommends migrating the workloads to these RIs. The plan assumes that an RI will always be less expensive than the on-demand counterpart, however this is not always the case. Because of some billing details from the service provider, the plan can recommend migrating to an RI that is more expensive than running on-demand.
- For Azure environments, the user interface can show storage entities for Standard SSD, even though these entities are not active in the environment. Turbonomic does not consider these entities in its analysis.

- To calculate whether to purchase RIs, either in real-time analysis or in Cloud plans, Turbonomic requires a full week of historical data to determine which VMs are good candidates for RI pricing. However, if you recently installed an evaluation copy of Turbonomic, it is likely that you have not gathered a full week of historical data.

For versions 6.3.1 and later, you can configure Turbonomic to recommend RI purchases with less than one week of data. To set up this configuration, contact your sales representative, or Technical Support.

Note that this configuration is for *simulating* RI Purchase decisions. If you install Turbonomic for a production deployment, you must be sure to reset the configuration to the default. After resetting the configuration, Turbonomic will use a full week of historical data when calculating RI Purchase actions.

- For Azure environments that configure Consistent Resizing for a scope, if you run a Migrate To Cloud plan to migrate the VMs to AWS, then the plan actions do not maintain consistent resizing for the VMs.
- For environments that make use of Price Adjustments or Azure Rate Cards, whenever you add, remove, or change a Price Adjustment or Rate Card that is in use, you must allow sufficient time for Turbonomic to fully discover all of the affected environment, and to propagate the changes throughout that environment. In an average environment, this can take up to 30 minutes. As an alternative, you can manually execute rediscovery for the affected cloud subscription or account.
- In AWS environments, for some templates the AWS API returns template values that do not match the AWS documentation. Specifically, the values for IO and Network bandwidth can differ from the values stated in the AWS documentation.

These values do not affect the actions that Turbonomic calculates for workloads on the cloud. However, for the affected templates the user interface can display bandwidth values that do not match the values you expect to see.

- When running Optimize Cloud plans, you should not add workloads to the plan. Adding a workload to this type of plan can give you inconsistent results.
- After completing an online upgrade, the User Interface continues to show the "Getting Ready" message. To remove the message and display the Login page, refresh the browser.
- In AWS environments, availability zones can support different sets of instance types (templates). Turbonomic can identify valid templates for a given *region*, but not for a given availability zone.

If Turbonomic recommends resizing to an invalid template, it could be that you have encountered this problem. To work around it, you can create a policy for the affected scope of workloads and manually exclude the invalid templates.

- For a Migrate to Cloud plan, when migrating from AWS to Azure, the plan does not place Spot Instances from the AWS environment.
- **Customer Issue 102202**

After you have enabled Datastore Browsing, you might want to disable it again. Under some circumstances, even though you change the setting and then restart the tomcat server, Turbonomic continues to execute datastore browsing.

If this situation arises, please contact Technical Support for assistance.

Also note, after you disable datastore browsing, it can take up to 10 minutes for the change to take effect, even after a tomcat restart.

- For the Optimize Cloud plan, as it calculates NFUs and other metrics Turbonomic introduces rounding errors. For smaller environments the error is on the order of 0.02. However, for larger environments the error can multiply into a significant range.

Affected plans are still useful. However, for a large planning scope the plan can show exaggerated RI Buy recommendations. To get more precision in your Optimize Cloud plans, run them on a smaller scope of workloads.

- **Customer Issue 100341**

In VMM environments, when a VM storage is on an SMB share and it also has an ISO image, then Turbonomic will not recommend storage moves from the SMB share for that VM.

- When running plans to migrate workload to the cloud, be sure to choose the **Migrate to Cloud** option, and do not use the **Workload Migration** option. If you use **Workload Migration**, you can successfully set up a migration to the cloud, but the resulting plan might not choose the least expensive regions for workload placement.

- **Customer Issue 100909,101615**

In some vCenter Server environments that include unusual configurations for vCenter, discovery of VMware vSphere Storage Policy Based Management can time out and fail to complete. If you encounter this situation, please contact Technical Support.

- **Customer Issue 100911,104082,104121**

For AWS environments that include Third Party Marketplace AMIs, Turbonomic can recommend scaling to an instance type that is not supported for a given Marketplace AMI. As a result, the scaling action can fail.

As part of its intelligent workload management, Turbonomic can recommend scaling the instance to a different instance type. However, Turbonomic chooses from the full set of available instance types, and does not recognize whether the recommended type is supported for the given Marketplace AMI. For this reason, the scaling action can fail.

To avoid this situation, create an Automation Policy for a group of related Third Party Marketplace instances, and then set the Excluded Templates scaling constraint to exclude any unsupported instance types.

- In AWS environments, under certain circumstances Turbonomic can show negative costs in charts for specific regions. This can occur when the AWS target indicates that it offers templates, but those templates are not truly available. In that case Turbonomic calculates negative values for costs in that region.
- In OpenStack environments, when you set up reservations to deploy workloads via OpenStack templates you must constrain the deployment to the OpenStack datacenter.

When you define the workload to deploy, you specify a template to deploy and any constraints that you want Turbonomic to respect. To deploy an OpenStack template, turn on the **Limit initial placement to locations you specify** constraint, and manually choose the datacenter or datacenters that support the OpenStack template.

- Turbonomic now requires HTTPS to connect to the user interface, and to connect to the API. Also, API connections no longer accept clear-text authentication, so you cannot include authentication in URLs to execute API commands. Instead, you should use curl commands to execute API commands.

To support HTTPS, Turbonomic includes a self-signed certificate by default. We recommend that you install a certificate from a trusted Certificate Authority. If you do not install a trusted certificate, you can still use curl to execute API commands if you include the `-k` flag.

- After editing the IP address of a Turbonomic target or deleting a Turbonomic target, we recommend that you restart Tomcat in order for the Supply Chain to correctly reflect the changes.

- Turbonomic policies include a default policy named **Global Actions Mode**. You can use this policy to globally disable all actions in Turbonomic. If you turn on the setting to disable all actions, then `Disabled` takes effect for all actions. However, in the default policies the action modes remain as you have set them. The user interface does not show that you have globally disabled these actions.
- For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.
- In AWS environments that use RI templates across access regions, you should be sure to have a single master account, and include that master and all sub accounts as Turbonomic targets.

If you experience situations where RI actions seem to recommend templates that are not available in the indicated regions, verify that you have included all the sub accounts as targets. If the problem persists, contact Technical Support.

- After executing **Settings : Maintenance Options : Export Current Environment**, the user interface sometimes shows that the export failed even though the export is still running successfully.

If you are exporting a large topology, the user interface response can time out and show this warning after 60 seconds, even though the export continues to run.

If you experience this situation, restart the Turbonomic server. This restarts the HTTP server with a 10 minute timeout setting, which should be sufficient to export most environments. If the problem persists after a restart, contact Technical Support.

- Storage Suspend actions appear grouped with Delete Wasted Files actions in the Delete category of pending actions.

For example, assume you are viewing the On-prem Overview on the Home Page, and the Pending Actions chart does not list Suspend Storage actions. However, if you hover on the Storage tier of the Supply Chain, the tooltip shows that you have actions to execute on Storage. These actions would be to suspend storage.

When viewing the overview, it is always a good idea to hover on any tier of the Supply Chain that is not completely green. The tooltip gives you extra information about that tier, including a count of actions.

To see the specific storage actions, click the Storage tier in the Supply Chain, and then view the Actions list.

- When running a Replace Workload plan to replace certain VMs with a template, under some circumstances the plan replaces the VM with a copy of an existing VM. For this reason, you should not configure a Replace Workload plan.

To get the same effect as a Replace Workload plan, you should configure the plan to remove the VMs you want to replace, and then configure it to add instances of a template or VM copy. For example, remove three VMs, and then add three instances of a VM template.

- For Load Balancer entities, the Transactions Per Second data can be incomplete.
- When setting up a custom plan, the user interface allows you to set a scope that includes public cloud entities. If you run such a custom plan, you can see unexpected or incorrect plan results. You should not set a scope that contains public cloud entities.

To run plans for the public cloud, choose the **Migrate to Public Cloud** or the **Optimize Cloud** plan type.

- In order to add an ACI Target, your Turbonomic instance must be using Market 1, which is disabled by default. If you require ACI Integration, please contact Turbonomic Technical Support.

- Turbonomic includes Automation Policies that you can create to modify the analysis of elements in your environment. This includes Scaling Constraints, which can include setting up templates to exclude when resizing workloads on the cloud. Note that the **Excluded Templates** option is not available for default policies for VM, Database, and Database Server templates.

To set up excluded templates, define a scope of entities and create an Automation Policy for that scope.

- **Customer Issue 99189,99805**

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.

Turbonomic considers storage latency when calculating whether to move a VM to existing storage, or whether to provision new storage. Because of this known problem, Turbonomic 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 Turbonomic 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`.
- 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.
 - Turbonomic can automate resize actions for datastores. However after executing the action, the hypervisor that is stitched to the datastore requires a refresh before it can recognize that change. Because Turbonomic uses the hypervisor to discover the datastore, then Turbonomic will not recognize the change either. As a result, Turbonomic might recommend the same action again, even though the datastore has already been resized.

To avoid this situation, Turbonomic suggests that you set the Action Mode for storage resize actions to be `Recommend`.

- **Customer Issue 98910**

In action scripts, you can get the internal name of a VM and use that to assemble calls to the API that work with the given VM. However, with the JSON API you must use the VM's UUID to access it via the API call to `entities`. Calls that use `VM_TARGET_NAME` with the XML API must now use `VM_TARGET_UUID`.

For example, the following code creates the URL base for a call to the API that will get the actions associated with the action script's target VM:

```
URL="/vmturbo/rest/entities/$VM_TARGET_UUID/actions"
```

- Turbonomic generates special average or max utilization templates that it uses when calculating cluster headroom. You should not edit these templates, because Turbonomic will overwrite your changes the next time it generates the templates. However, the Template Catalog presents these templates as editable.

- In AWS environments, Turbonomic discovers data from a given AMI to determine the OS of an associated VM. The OS influences cost calculations.

If you delete the AMI from your environment, then Turbonomic cannot discover the OS, and so it assumes a free Linux OS. This can result in incorrect calculations of cost in plans and in evaluations of real-time costs.

- After restarting the Turbonomic server, users must log into new sessions in order to continue using the user interface or the API.
- For Migrate to Cloud plans, when you migrate a VM that has an attached ISO image, the plan shows a move of a non-existent OGB disk. This OGB disk is a representation of the ISO image, and you can ignore the move action. The other plan actions for the VM are correct.
- For Tomcat, SQLServer, WebSphere, and other application or database targets that use a scope to identify target instances, Turbonomic can fail to validate or discover the targets. If you add a target via scope, and that scope does not have any VMs to host the target applications, then the target will not validate. If you later add hosts for the applications to that scope, Turbonomic does not dynamically recognize the change and then validate and discover. Even if you execute a Validate command for that target, Turbonomic can validate but it will not run discovery.

To avoid this problem, make sure your applications are running on hosts before you configure the target. If you have encountered this problem (you added hosts to a scope after configuring the target), delete the target from Turbonomic and set a new target with this scope.

- When you set up action orchestration in an Automation Policy, you should be sure that the scope for this policy does not include conflicts with individual entities. If a single entity is in two scopes, and one scope enables orchestration while the other scope does not, then Turbonomic arbitrarily chooses either orchestration or a Turbonomic action mode.

For example, assume two groups - GroupA and GroupB, and assume a host named MyHost is in both groups. If you configure action orchestration for hosts in GroupA but you do not for hosts in GroupB, then you cannot be sure that Turbonomic will call the action orchestrator for actions on MyHost.

- Assume you have application or database servers as targets, and they use dynamic groups to define their scopes for monitoring. If you add new application or database servers to these dynamic groups, or if you shut down and then restart an existing server, then Turbonomic fails to discover the change and these servers will not appear in the user interface. To resolve this problem, execute a manual rediscovery of the affected target.
- Under some circumstances when using the Deploy View, the user interface can fail to respond. After you request a placement, if the placement recommendation does not appear within one to two minutes, reload the browser running Turbonomic to reset the user interface.
- When you reserve resources for VMs that you will deploy in the future, Turbonomic initially calculates these reserved resources as CPU Provisioned, Memory Provisioned, and Storage Provisioned. However, these resources ultimately revert back to values that do not include the reserved VMs.
- There is a rare case where Turbonomic can lose the cluster relationship for VMs running on a host. Assume you move a host out of a cluster, and directly into the datacenter (not into a cluster). Next you create a new cluster and then move the host into that cluster. In this case, the VMs on that host will not belong to any PM cluster. This can also affect Provider Virtual Datacenters that belong to the new cluster.

To avoid this problem, create the cluster first, and move the host directly into it. If this problem occurs, rediscover your environment to establish the correct cluster relationships.

- For VMs running on Hyper-V, if you set a VCPU limit (limit VCPU to less than 100%), then the VCPU utilization data that VM returns to Turbonomic is not correct. As a result, Turbonomic will not recommend that you increase the VCPU limit.
- In OpenStack environments, it is possible to place a VM belonging to a specific cluster (a VM flavor that is set via extra specifications) onto a host that is not a member of that cluster. Turbonomic cannot identify this configuration error, and will not recommend a move to place the VM on an appropriate host.

You can implement placement policies to ensure that VMs always get placed on the correct hosts. This can work even if there is no host cluster to match the VM flavor. However, in that case the user interface will not show these VMs as members of a PM cluster.