



Turbonomic 6.3.0 Release Notes

February 19, 2019

This document describes issues that are addressed in Turbonomic 6.3.0 – Release Date: February 19, 2019. All builds are cumulative. Applying 6.3.0 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>

What's New for Version 6.3

The 6.3 family of Turbonomic releases includes the following new features:

- Price Adjustment for Cloud Environments

This release improves how we handle price adjustments. In previous versions, you could configure discounts for specific billing groups in your cloud environment. This release extends that feature to enable more general price *adjustments*, where you can specify a discount or an increase in cost. Further, you can drill into line items and override the overall adjustment as necessary.

See "Price Adjustments" in the *User Guide*.

- Azure Rate Cards

For Azure environments, Turbonomic uses the published Microsoft rates to calculate costs for templates and services in the cloud. However, your environment can run workloads on an Azure environment that is managed by a specific Cloud Solution Provider (CSP). That CSP can offer its own price list for the Azure services. In this case, you can upload the CSP price list to create a Rate Card, and configure Turbonomic to use those rates for the affected Azure subscriptions.

See "Azure Rate Cards" in the *User Guide*.

■ Cloud Planning Improvements

This release introduces the following features to enhance the value of running plans in your cloud environment.

- Plan Migrations to Azure Cloud Solution Providers

This release introduces Azure Rate Cards. With these rate cards, when you plan a migration to a CSP that uses its own price list, the plan calculates the true cost of that migration. In this way, you can use a Migrate to Cloud plan to illustrate the value of moving to a subscription hosted by the given CSP.

See "Migrating to Azure CSPs" in the *User Guide*.

- Plan migration of Physical Hosts to the Cloud

With the introduction of the BareMetal target, you can configure a set of physical hosts for Turbonomic to discover, and treat as VMs. Then you can run a plan that migrates those VMs to the cloud.

See "Bare Metal Hypervisor" in the *Target Configuration Guide*.

- Include Azure RI Purchases in Migrate to Cloud Plans

Turbonomic can now recommend Azure RI purchases in migration plans.

- Consider Future RI Purchases in Optimize Cloud Plans

When you run the Optimize Cloud plan, you can now set a purchase date for RIs. This date reflects when you intend to execute the recommendations that the plan calculates. As the plan calculates RI purchases, it will not include RIs that would be expired before that purchase date.

See "Planning Cloud Optimization" in the *User Guide*.

- Scope to Billing Family in Optimize Cloud Plans

With this release you can limit an optimization plan to a single AWS billing family. The plan scope includes all the accounts in that family, and the plan calculates RI purchases through the billing family's master account.

■ New analysis policies for management on the public cloud

This release includes new policy settings to improve the way Turbonomic calculates actions for workloads on the public cloud. For more information, see "Analysis Policies: VMs" in the *User Guide*.

- Consistent Resizing

When you turn on Consistent Resizing for a group of VMs, Turbonomic resizes all the group members to the same size, such that they all support the top utilization of each resource commodity in the group. In addition, Turbonomic discovers the members of Azure Availability Sets and AWS Autoscaling groups, and automatically enables Consistent Resizing for these groups.

- Instance Store Aware Scaling

For AWS environments, the template for your workload determines whether the workload can use an *instance store*, and it determines the instance store capacity. With this setting, you can ensure that resize actions respect the instance store requirements for your workloads.

■ Set Turbonomic Scope to AWS Billing Families

As Turbonomic discovers AWS billing families, it includes them as a type of group in the **Search** page. You can navigate to **Search**, choose **Billing Families**, and then set the Turbonomic scope to one or more billing families.

■ Improved Discovery of Database Information

For AppDynamics and MS SQL targets, Turbonomic now discovers the database type (MS SQL, MySql, Oracle), and when possible it discovers the database version as it discovers database entities. You can filter by this information when creating groups, and Turbonomic displays this information in the user interface.

■ Improved Charts in the User Interface

• Improvements for the Cloud Cost Comparison Chart

Turbonomic analyzes your cloud environment to calculate the best placement and scaling of workloads, and to recommend the actions to optimize your environment. The Cloud Cost Comparison chart compares your current costs with the costs you would see if you execute the pending actions. We have redesigned the chart to make it easier to read, and to display more relevant information.

You can see this chart in the Cloud View of the Home Page, the Optimize Cloud plan results, and the Migrate to Cloud plan results.

• Display Cost in RI Utilization Charts

RI Utilization charts show utilization of your RI inventory in terms of Normalized Factor Units (nfu). With this release the charts include an option to display utilization in terms of dollar amount.

■ Workload Resize and Move Improvements on the Cloud

• Workload Processor Compatibility

Turbonomic recognizes processor types that you currently use for your workloads. For move or resize actions, Turbonomic keeps your workloads on instance types with compatible processors. For more information, see "Azure Template Requirements" in the *User Guide* and "AWS Template Requirements" in the *User Guide*.

• Recovery for Failed Resize and Move Actions

For workload on the public cloud, if Turbonomic tries to execute a move or a resize action but the action fails, then Turbonomic places the affected VM in a special group named *Cloud VMs with Failed Sizing*. You can review the group, and inspect individual VMs to troubleshoot the cause of the failure. As soon as Turbonomic successfully executes a move or resize on a VM in this group, it then removes the

VM from the group. For more information, see "Cloud VMs with Failed Sizing" in the *User Guide*.

■ **More Supported Versions of Hypervisor Platforms and Targets**

With this release, Turbonomic supports installation on VMware vCenter Server 6.7.

In addition, this release adds target support for Microsoft SQL Server 2017.

■ **New Targets**

This release of Turbonomic introduces the following new targets:

• **Dynatrace Server**

Turbonomic now supports discovery of applications that are managed by the Dynatrace platform. Turbonomic includes the discovered information about these applications in its calculations for VM actions.

See "Dynatrace" in the *Target Configuration Guide*.

• **Bare Metal Hypervisor**

You can configure Turbonomic to discover specific physical hosts that you have in your on-prem environment. When Turbonomic discovers them, it treats them as special host-proxy VMs. Turbonomic cannot execute any real-time actions on these host proxies. But you can run plans to migrate these proxies onto the public cloud.

See "Bare Metal Hypervisor" in the *Target Configuration Guide*.

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.

Update Links

Turbonomic 6.3.0 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.

Security Requirements for Browsers

For web browsers to communicate with Apache, the Apache configuration requires TLS version 1.1 or later. To use versions of Microsoft Internet Explorer 9 and 10, you must enable TLS 1.1 or later (in **Internet Options > Advanced**). For more information, see the following Green Circle article:

[The DROWN Attack: Configuring your Turbonomic Web Security](#)

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`

Configure Email Notifications for Database Disk Usage

You can configure Turbonomic to notify you by email whenever the database is down or storage exceeds 80% utilization. To configure this notification, you will first execute a script

that initializes the configuration. Then to set up a regular check of disk consumption, you will set up a cron job to execute the scripts Turbonomic provides on its server.

To perform the one-time configuration for this notification:

1. Execute one of the notification scripts.

To initialize the configuration, you can execute either of the provided scripts. To do so, execute one of the following commands:

- `/srv/tomcat/script/appliance/turbo_check_db.sh`
- `/srv/tomcat/script/appliance/turbo_check_disk.sh`

2. Provide the information as the script prompts you for it.

The script displays the following prompts:

```
[root@turbonomic appliance] ./turbo_check_db.sh
Configuration file does not exist.
Creating default configuration file.

Database name (Default: vmtdb):
Database user (Default: root):
Database password:
Database port (Default: 3306):
Email address to send notifications to:
```

Be sure to give the email address where you want to receive the notifications.

Running the script creates a configuration file that saves the settings you provide.

To set up regular checks that run every 30 minutes, append the following lines to your crontab file:

```
*/30 * * * * /srv/tomcat/script/appliance/turbo_check_db.sh >/dev/null
2>&1
*/30 * * * * /srv/tomcat/script/appliance/turbo_check_disk.sh >/dev/
null 2>&1
```

After you set up the cron job, if the database goes down, or if disk usage exceeds 80%, the scripts send an email alert to the recipient that you specified in the initial configuration. In addition, the scripts write alerts to the log file: `/var/log/tomcat/monitor.log`.

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.proxyPassword

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.proxyPassword=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 104393**

Fixed: Assume you add an Application Server target that accepts multiple IP addresses (for example, Tomcat Server or MySQL Server), and you provide multiple IP addresses for the target. If the first connection via an IP address fails, Turbonomic does not attempt to connect via the other IP addresses you provided, and the validation fails.

- **Customer Issue 104417**

Fixed: When you run a plan to add workloads in a scope that includes suspended hosts, the plan results can be incorrect. For example, the plan can provision new hosts even though that action is disabled for the plan, or it can place workloads such that they consume more than 100% of host resources.

- **Customer Issue 103587**

Fixed: When you create VM snapshots while backing up your VMs, Turbonomic can recommend a Storage Move to accommodate the surge in demand. In some cases, the target storage does not have enough capacity, so the move fails.

■ **Customer Issue 104080**

Fixed: In Azure environments where you have added a discount to some subscriptions, if the scope includes storage entities that are not included in the discount, then the user interface can show inconsistent costs in different charts on the Cloud Storage Devices page. The Potential Savings chart applies the discount to the storage, while the Estimated Cost Breakdown by Unattached Tier chart does not.

■ **Customer Issue 104083**

Fixed: In AWS environments, Turbonomic can recommend resizing a workload that uses encrypted storage, to a template that does not support encrypted volumes.

■ **Customer Issue 103989**

Fixed: When VMs in an Azure environment shut down, the Azure account continues to incur storage costs, but the Cloud Cost Comparison chart shows the cost as 0 (zero).

- **Fixed:** For Migrate to Cloud or Cloud Optimization plans in a very large environment, under some circumstances the plan performance is unacceptable because of an internal scoping issue.

■ **Customer Issue 103853**

Fixed: In public cloud environments that have no pending RI Purchase actions, the Cloud Cost Comparison chart on the Home Page should display values similar to the same chart for an Optimize Cloud plan that has run over the same scope. However, under some circumstances these charts can differ significantly.

■ **Customer Issue 103853**

Fixed Under some circumstances, the RI Utilization chart can show overall utilization exceeding the actual RI inventory.

■ **Customer Issue 103454**

Has workaround: Reports fail to generate when you change the Turbonomic port from 443 to some other valid port number. Contact Technical Support for the workaround to this issue.

■ **Customer Issue 103465**

Fixed: In Azure environments that configure quotas for VCPU resources, Turbonomic should not recommend resize actions that will exceed the configured quotas. However, under some circumstances the quota information that Turbonomic stores can be out of sync, and so it can recommend actions that exceed quotas.

■ **Customer Issue 103466**

Fixed: In Azure environments, resize actions do not respect the instance store requirements for workloads. As a result, Turbonomic can recommend moving a

workload to a template that does not support Instance Stores, or it can move a workload to a template that provides less capacity or a lower tier for the instance store.

■ **Customer Issue 103487**

Fixed: After running a Migrate to Public Cloud plan, the plan results show incorrect storage sizes for the current on-prem storage.

- **Fixed:** In AWS environments, under some circumstances an Optimize Cloud plan can halt with an exception, which produces incomplete plan results.

■ **Customer Issue 103341**

Fixed: When configuring a plan you can set an increase to the overall utilization of resources in the plan scope. However, this setting does not take effect when you run the plan.

■ **Customer Issue 102985**

Fixed: Under some circumstances, the Pending Actions list shows duplicate recommended actions to provision a specific host. The user can manually execute some of the actions, even though the action mode configured in the host policy is Recommend.

- **Fixed:** In the Multiple Resources chart, the custom time range set by the user is not correctly reflected in the chart.

■ **Customer Issue 101724,102154,102546,104187,104469**

Fixed: Under some conditions, when Turbonomic displays a resize action for a workload, the RESIZE FROM and RESIZE TO show values of 0.0.

■ **Customer Issue 102236**

Fixed: In Azure environments that support Accelerated Networking, Turbonomic can recommend moving a workload that uses Accelerated Networking to a template that does not support it.

- **Fixed:** In Azure environments, Turbonomic can discover VM groups based on tags. However, if the tags use the square bracket character, then Turbonomic does not create the group. For example, Turbonomic does not create a group for the tag, `test:[MyTag]`.

■ **Customer Issue 101954,102057,102462,102947,103017,103387**

Fixed: In large environments, the performance for loading the On-prem Executive Dashboard is unacceptable.

- **Fixed:** The user interface to edit charts is inconsistent, and disables settings that should be enabled.
- **Fixed:** For AWS environments, the template for your workload determines whether the workload can use an instance store, and it determines the instance store capacity. As Turbonomic calculates a resize or move action, it can recommend a new template

that does not support instance stores, or that does not provide the same instance store capacity.

■ **Customer Issue 102691,103954**

Fixed: For Azure environments, Turbonomic does not provide a way to identify Availability Sets and ensure that all the workloads in the set will resize consistently to the same template.

- **Fixed:** For WebSphere, Tomcat, and SQLServer targets, you can set a scaling policy to be horizontal or vertical. To set scaling for these targets, you must set it both in the Turbonomic user interface, and also in the properties file for the given probe. This means that for any of these targets, all discovered entities must use the same scaling property. For help with scaling policies for these targets, contact Technical Support.

Known Issues

- For Azure environments, Turbonomic does not support workloads running in Early Updates Access Program (EUAP) regions. If you have a workload running in such a region, Turbonomic will try to discover that workload, but discovery will fail for the affected Azure target.

- For Optimize Cloud plans, you should not use the **Purchase Date** setting for RIs when you configure the plan's RI Purchase Profile.

The Purchase Date setting sets when you intend to execute the optimizations that the plan calculates. When you set a purchase date, the plan should not recommend any RIs that expire before that date. However, this feature is still experimental, and the plan might not show the correct results for future RI purchases.

- Starting with version 6.3.0, you can no longer specify **Target RI Coverage** in policies and plans. Earlier versions of Turbonomic include a control to set a specific target percentage for RI coverage for a scope of your environment. You could make this setting for real-time analysis and for Optimize Cloud plans. With this version you can no longer make the **Target RI Coverage** setting. However, the plan results page shows an entry for RI Coverage Override. That setting has no effect.
- In the user interface, you can scope the view to a single VM. In that case, the charts show details about the given VM. When you set the view time frame to 2 Hours, then some of the charts show incomplete data, or no data. If the data you need to see is missing or incomplete, set the time frame to 24 Hours or greater.
- In the Cloud Cost Comparison chart, if you set the scope to include part of a billing family but not all of the billing family, the value for Reserved Compute Cost shows the cost for the entire billing family, not for just the RIs in the current scope.
- Starting with version 6.3.0, you can no longer specify **Target RI Coverage**. Earlier versions of Turbonomic include a control to set a specific target percentage for RI coverage for a scope of your environment. You could make this setting for real-time
-

analysis and for Optimize Cloud plans. With this version you can no longer make the **Target RI Coverage** setting.

- Cost Breakdown by Component charts show the current breakdown for costs, and also show projected costs after executing pending actions. If the pending actions include RI Purchase actions, the projected costs can be incorrect. The current and historical costs are correct.
- In Optimize Cloud plans, the Cloud Cost Comparison chart includes entries for Current costs as a way to compare your cost today with your cost after running the plan. If you specify an expiration date for the RIs in your plan, then these "Current" entries show the cost as of that expiration date. In that sense, the label "Current" is misleading.
- 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 in the service provider, the plan can recommend migrating to an RI that is more expensive than running on-demand.
- In the user interface, charts for Application Guestload show incorrect values for VCPU and VMEM.
- In Migrate to Cloud plan results, the Cloud Cost Comparison chart sometimes displays an inaccurate value for the Workload with Efficiency Opportunities row.
- This version of Turbonomic includes changes to the Cloud View on the Home Page. For example, the Cloud View no longer includes the Cloud Cost Comparison line chart. However, after you update to this version the Home Page continues to show the old view, and the Cloud Cost Comparison line chart can show stale or incorrect values.
After updating to this version, you should click **Reset View** (at the top-right of the home page) to display the default layout for the Cloud View.
- In Azure environments that include Availability Sets, if a VM in an Availability Set uses Accelerated Networking, then Turbonomic can recommend resizing the VM to a template that does not support Accelerated Networking.
- For AWS 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.
- In Pure Storage environments, if you have more than 500 volumes managed by a single storage controller, then Turbonomic does not collect the volume statistics for that storage controller. Note that this is an unusual situation, and 500 volumes exceeds the recommended count for a single storage controller.

■ **Customer Issue 103853**

Fixed - Needs Restart: Performance for the user interface is unacceptable because the browser does not cache images and other assets.

To enable this fix, you must restart the Apache HTTPD server after you upgrade to the latest version.

-
- 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.
 - The Policy Settings for VMs include a setting named **Enable High Availability**. The documentation states that this setting sets the minimum number of VMs to place in a given scope. However, this setting has no effect.

■ **Customer Issue 103513**

For Storage policies, you can enable datastore browsing, and then you can specify which files or directories to ignore via the **Files to Ignore** and **Directories to Ignore** settings. However, these parameters only take effect if you set them in the global Storage policy.

- The user interface includes Action Policy settings that Turbonomic analysis does not support. These settings are:
 - VM Action Policy: OperationalConstraints/Storage Latency SLA Value
 - Application Action Policy: Application Server Discovery/SLA Capacity

■ **Customer Issue 103310,103733**

If you have configured your Turbonomic to authenticate users via Single Sign-On (SSO), then you should be careful to not let your Turbonomic license expire. Because SSO currently works as a licensed feature, if your license expires then you can no longer log in via SSO.

If your license does expire and you cannot log in via SSO, perform these steps to recover SSO capabilities:

- Disable SSO for your installation
For instructions, see *Disabling Single Sign-On* in the Installation Guide.
- Add or upgrade your Turbonomic license
Log in as an administrator user via a local user account, and then add or upgrade your license.
- Enable SSO again
- For instructions, see *Configuring Single Sign-On* in the Installation Guide.

- 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 Azure environments, Turbonomic discovers Resource Groups and presents them in the user interface as groups. It also discovers Cloud Solution Provider business accounts, and Azure subscriptions.
-

One thing you can do is to create your own groups of these discovered resource groups, business accounts, or subscriptions. When you upgrade from a version earlier than 6.3, the upgrade cannot retain these types of groups that you have created. If you had created groups of resource groups, business accounts, or subscriptions in the earlier version, then you must create them again after the upgrade.

- As you update Turbonomic from a version that is earlier than 6.3, it is possible for the upgrade to log errors that contain statements similar to `ERROR [main] [PoliciesConfigTopologyMigrationScript] : Setting invalid boolean value ActionScriptManager attr terminatePM to 'false';`. This can occur as the upgrade encounters invalid values in your Turbonomic policies.

Earlier versions of Turbonomic silently ignored these errors, and used the default value whenever it encountered an invalid setting. Starting with version 6.3, Turbonomic logs the errors, and sets the affected value to its default.

If you encounter these errors, you should review your policies to make sure they are as you intend. You can look at the log messages to see the type of setting that is affected. In the above example, you can see that a Terminate PM setting was invalid, and it is now set to `false`.

- While configuring or running Migrate to Cloud or Optimize Cloud plans, you can encounter various issues with the user interface:
 - Using **Advanced Configuration**

Plan configuration includes **Advanced Configuration**, which sets up whether to include RI purchases, and which purchase profile to use. If you make these advanced settings and then navigate to make other plan configurations, your advanced settings will be lost. You should make the advanced settings as the last step of your configuration, and execute the plan from that screen.
 - Re-running a Migrate to Cloud Plan After Turbonomic Restart

In this circumstance, the plan scope includes the workloads that had already been migrated in the first run of the plan. After a restart, be sure to check the plan scope before re-running a migration plan.
 - Re-running a Migration Plan After Adding Workloads to the On-Prem Environment

After you run a migration plan, assume you have added workload to the on-prem environment. If you then re-run the plan, its scope will include the new on-prem workloads. Be sure to check the plan scope before re-running a migration plan.
- For Azure environments, there is a known issue with access to the data for VMs in the Germany regions. Specifically, the environment does not reliably return pricing information for Managed Standard SSD disks. As a result, Turbonomic might not display accurate costs for workloads in these regions.
- In AWS environments, for some templates the AWS API return 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.
- In vCenter environments, to support VM moves across vCenter Server boundaries, the documentation instructs you to create merge policies to merge host clusters and to merge storage clusters. This enables Turbonomic to treat the merged clusters as one pool that it can move VMs into. In this way, you can move VMs from one datacenter into another.

However, the documentation does not tell you to also create a merge policy for the affected datacenters. If you do not create this datacenter merge policy, then after executing moves across datacenters, Turbonomic will post Reconfigure actions for the affected VMs.

Before enabling cross-vCenter moves, be sure to create the necessary merge policies for the affected datacenters.

If you already enabled cross-vCenter moves and then encounter this problem of Reconfigure actions:

- First create the necessary merge policies for the affected datacenters
- Then restart your Turbonomic

- When you configure a Cisco HyperFlex target, you can specify target credentials for the root user. For Turbonomic versions 6.2.1 and earlier, the root username must be set as `root`. For later versions, the username must be set as `local/root`.

Because of this change, as you upgrade to a later version, HyperFlex targets with a root account will fail to validate until you edit the target to specify `local/root`.

- 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.
- 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.

- In AWS environments, when you scope to a region then Turbonomic does not show any pending cloud-based actions. To see all of your cloud actions, set the scope to something other than an AWS region.

■ 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.

■ Customer Issue 101794

In the classic version of Turbonomic the user interface displays an audit log that lists all the actions that have been taken in your environment.

In earlier versions, for vCenter environments this log included actions that were taken by the vCenter Server. For example, if a user used vCenter to move a VM, or if DRS rules made a change, the audit log would show those actions.

Starting with version 6.1, the audit log no longer shows these external actions.

- 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.

- When using the API to deploy reservations (using Deployment Profiles and Reservations endpoints), the placement calculations can fail to respect network constraints. This occurs when one or more networks have the same name. The API can fail to recognize the network constraint, even if you identify the affected networks via `UUID` values.
 - 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 `$VMT_TARGET_NAME` with the XML API must now use `$VMT_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/$VMT_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.

■ Customer Issue 98320

Customer Issue 104461,98320

In Turbonomic you can create policies with a scope to a dynamic group. If changes in your environment remove all the members of the group for that policy, then Turbonomic deletes the policy definition.

-
- To ensure that Turbonomic recommends the most appropriate actions for AWS environments, you must make specific settings for default policies in your Turbonomic installation. After adding an AWS target, then perform these steps:

- Click **Settings: Policies** to navigate to the Policy Management page.
- Ensure datastore browsing is enabled for the Storage Defaults policy.

Find the Storage Defaults policy and make sure that the option for Disable Datastore Browsing is OFF.

- 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 0GB disk. This 0GB 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.
- Assume you set up a scope that is less than a full datacenter. If you ignore hyperthreading in that scope and then restart tomcat for the Turbonomic application, then Turbonomic does not calculate the CPU capacity correctly for the affected scope of hosts.

If this problem occurs, perform a full rediscovery to show correct CPU capacity.

- 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.