



Turbonomic 6.3.7 Release Notes

May 31, 2019

This document describes issues that are addressed in Turbonomic 6.3.7 – Release Date: May 31, 2019. All builds are cumulative. Applying 6.3.7 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 6.3.7 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

- **Fixed:** For public cloud environments, if you create a static group of accounts and those accounts do not contain workloads of a given type, then as you add workloads of that type they do not appear in the group. For example, if your accounts do not contain any Database Server instances when you create the group, if you later add a Database Server instance to an account in the group, that Database Server will not appear within the scope of the static group.
- **Customer Issue 105732,106025,106226**
Fixed: Plans to add workload fail to add the indicated workload to the plan market.
- **Customer Issue 106029**
Fixed: For Azure environments, under some circumstances Turbonomic does not discover all of the metric data for the inventory. There are cases where Turbonomic does not recognize the names of metric groups that the target returns.
- **Fixed:** For AWS environments that include Billing Families, if you have a master account for a billing family that does not include any workloads of its own (all the workloads are managed by other members of the billing family), then the Top Accounts chart does not display an entry for this master account.
- **Customer Issue 105957**
Fixed: Under some circumstances when a connection to a vCenter target fails, discovery halts. There are circumstances where the network configuration results in the following vCenter Server error: `The destination host did not receive data from the source host on the vMotion network (2145068)`. If that error occurs, then this failure of discovery can have unexpected results.
- **Fixed:** Under some circumstances the discovery of targets is delayed up to three hours.
- **Customer Issue 105878**
Fixed: In public cloud environments, scaling actions for some VMs show an action mode of RECOMMENDED when the action policy sets the mode to MANUAL. This can happen with VMs that have no RI coverage when they are resized to an instance that has full RI coverage.
- **Customer Issue 105929**
IMPROVEMENT: For AWS environments, this release improves the steps to configure CloudWatch Memory Metrics on Windows instances. For more information, see the Green Circle article, [Setting up CloudWatch Memory Metrics on Windows Instances on AWS](#).

- **Customer Issue 105904**

Fixed: When you set up an email notification for Turbonomic events, the action to send the email fails, and you do not get any email notifications.

- **Customer Issue 105901**

Fixed: In a scope where you have set the action mode to DISABLED for Suspend Host actions, Turbonomic can still recommend actions to clear workloads off of hosts as though it is preparing to suspend the host.

- **Fixed:** After you update to version 6.3.7, any RI Coverage chart for a scope related to Business Accounts will not display historical data from before that update.

- **Customer Issue 105915**

Fixed: For public cloud environments, under some circumstances the user interface display of RI Utilization for a scope of workloads can exceed 100%. This can occur when updated calculations do not clear previously calculated data.

- **Customer Issue 105825**

Fixed: In public cloud environments, under some circumstances action details can show the incorrect costs for on-demand pricing. This can happen when there is some RI coverage for the affected workload.

- **Customer Issue 105737,106194**

Fixed: In AWS environments that have billing family accounts but no master account for the billing family, Turbonomic logs excessive messages for the error `There is no master business account.`

- **Customer Issue 105836,105896**

Fixed: For AWS environments with VMs that use EBS volumes, Turbonomic identifies active EBS volumes as inactive, and recommends removing those volumes. This is due to an error parsing the data returned by the target.

- **Customer Issue 105512**

Improvement: For discovery of Azure environments, discovery error messages include more information to describe the issue.

- **Customer Issue 105821**

Fixed: Under rare circumstances for public cloud environments, the tracking of cloud cost in the Cost Breakdown chart can show an anomalous spike. See also [Fixed: Duplicate Issue \(page 9\)](#).

- **Customer Issue 105820**

Fixed: The user interface uses the Pending Actions chart, and the Necessary Investments and Potential Savings charts to show the effect recommended actions have on cost. Under circumstances where actions do not result in savings *or* cost, these charts can show conflicting numbers of actions. This can happen for actions that resize a workload from one RI to another, where RI utilization remains the same (for example, from 100% to 100%).

To fix this issue, Turbonomic tracks actions with zero cost/zero savings in the Necessary Investments chart. With this change, both sets of charts show the same numbers of actions.

- **Improvement:** This release improves the performance of reservation and placement calculations, and improves the handling of concurrent reservation requests.

Placement calculations use the Turbonomic plan process to calculate the best placement for your workloads. The process can take more or less time, depending on the scope of your placement request, and the number of workloads you want to place. You can create new placement requests while Turbonomic is still calculating an earlier request.

If you request a reservation and the planning cannot place one or more VMs, then Turbonomic treats that reservation as *unfulfilled*. Turbonomic will continue trying to place that reservation until the request times out. By default, the timeout period is 120 minutes, and the retry interval is every 20 minutes.
- **Customer Issue 105678**

Fixed: Under rare circumstances for public cloud environments, the tracking of cloud cost in the Cost Breakdown chart can show an anomalous spike. See also [Fixed: Duplicate Issue \(page 8\)](#).
- **Customer Issue 105578**

Fixed: API requests for Stats over time do not include historical data for `numVCPU`s. These requests only return the current data for that resource.
- **Customer Issue 105465**

Fixed: When using a regex filter to create a group of AWS accounts, the filter fails to match all of the appropriate account names.
- **Customer Issue 105498**

Fixed: In environments with multiple VMAX targets, under some circumstances discovery can encounter duplicate storage IDs. As a result, Turbonomic does not include the affected storage in the topology that it manages.
- **Customer Issue 105378**

Fixed: In custom dashboards, if you add multiple area charts then the display of the data can be distorted.
- **Customer Issue 105469**

Fixed: For Azure environments, under some circumstances Turbonomic does not recommend scaling actions for VMs that should see such actions. This can occur when discovery of datastores is incomplete.
- **Customer Issue 104570,104585,105467**

Fixed: When you configure email notifications for issues in your environment, for some issues Turbonomic continually sends repeated notifications.

Known Issues

- You can create policies that exclude specific templates for a given scope. For example, in a public cloud environment you can choose a group of VMs and set a policy that excludes t2.small templates. In that case, Turbonomic will not recommend any actions that would scale a VM to the t2.small template.

For some groups of VMs, policies to exclude templates do not take effect, and Turbonomic recommends the excluded templates anyway. For example, if you apply the policy to a cloud account, or to a large group of VMs, template exclusion will not work. If this occurs, you can try using smaller groups as policy scope.

- **Customer Issue 105840**

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 106009**

When you create an Optimize Cloud plan, you should always configure the plan via the Optimize Cloud wizard. If you start the plan from a session scope, it is possible that you can include scopes that are invalid for this type of plan. Also, if you are planning for scopes that include billing families, it is possible that you can see inconsistent results.

For this reason, with Optimize Cloud plans you should always navigate to the Plan view, select Optimize Cloud, and specify the plan scope through that wizard.

- Because of changes to the SELinux configuration, you might need to temporarily disable SELinux on the Turbonomic server before you can execute a product update. If you have recently updated your server OS via `yum update`, then updates to Turbonomic can fail with the following error:

```
One of the configured repositories failed (Unknown)
```

To install the Turbonomic update in this case:

1. Disable SELinux for the Turbonomic server.

Open an SSH terminal session to the server and execute the command, `setenforce 0`.

2. Execute the offline update.

Follow the instructions for the update (see the *Installation Guide*).

3. Once the update is complete, enable SELinux.

In your terminal session, execute the command, `setenforce 1`.

After you successfully update Turbonomic, you will not need to perform these steps for subsequent updates.

- 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.

- The supported URL structure for API calls takes the form, `https://{Your_Turbonomic_IP}/api/v2/{resource_name}`. However, the Swagger-UI generates deprecated URLs when it shows the links in a response and when it shows the request URLs. These deprecated URLs take the form, `https://{Your_Turbonomic_IP}/vmturbo/rest/{resource_name}`.

As you implement scripts for production use, you should be sure to use the supported URL structure.

- Customer Issue 105495**

For schedule windows in policies, there are circumstances that can cause the schedules to fail when your system makes the change between daylight savings time and standard time.

Assume you create two or more policies that are scheduled back-to-back (start time of the next policy is less than an hour away from the end time of the previous policy) on the days for the daylight savings time change. In this case, policy creation can fail with an `overlapping schedule` error. In order to avoid this error, either exclude the daylight saving date or move the schedule for one of the policies by 1 hour.

- For Azure environments, when you deploy an instance of Turbonomic to Azure, it is possible that the initial boot can time out. As a result, the initial boot of the VM fails. If this occurs, try to boot the VM again, and log in as usual. As of this writing, this has only been observed for the initial boot. If the problem persists, please contact Technical Support.

- 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.

- For public cloud environments, Turbonomic does not support specialized pricing for dedicated instances or dedicated host tenancy. If it includes these instances in its calculations, then the pricing will be incorrect. This causes the resulting placement, scaling, or RI recommendations to be incorrect. For this reason, Turbonomic does not discover and represent dedicated instances in its topology.

- Customer Issue 104680**

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.

- To improve the performance for updates to version 6.3.3 or later, the update script does not replicate the associated risks for historical actions. After the update, any new actions will include the correct associated risk information, but previously existing actions will show a value of `unknown` for their associated risk information.

This occurs for upgrades *from* a version prior to 6.3, *to* version 6.3.3 or later.

If you require the historical associated risks, please contact Technical Support.

- **Customer Issue 104741**

In Hyper-V environments, it's possible that your environment uses Server Message Block (SMB) storage. You should be aware that Turbonomic Hyper-V discovery does not support SMB storage.

To manage SMB storage, Turbonomic requires a VMM target, and that VMM instance must manage the Hyper-V hypervisors and the SMB storage that they use.

If you manage a Hyper-V plus SMB environment via Hyper-V targets, then the data Turbonomic collects for SMB will be incorrect.

- **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 104698**

In Azure environments you can create disk storage and blob storage. The Wasted Storage report can show blobs as wasted storage, and you can read that to mean you should delete the blob to save costs. However, listing a blob as wasted storage is ambiguous, and in some cases a VM can actually use that storage.

When you run a Wasted Storage report, review it carefully before deleting any stores. In the case of blob storage, be sure that storage is not attached or in use by any of your workloads.

- **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.

- 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, as well as 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 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.
- In the user interface, charts for Application Guestload show incorrect values for VCPU and VMEM.
- 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.
- 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.

- **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 from an earlier version to version 6.3 or later.

- 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 103454**

Reports fail to generate when you change the Turbonomic port from 443 to some other valid port number. If you want to enable reporting after you set a custom port, contact Technical Support.

- 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
- 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. Advanced settings should be the last step of configuring your plan, and you should 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 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.
- 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.

- **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.

- 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.

- 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 Datastore Browsing is ON.
- 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.

- 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.