



Turbonomic 6.1.0 Release Notes

May 2, 2018

This document describes issues that are addressed in Turbonomic 6.1.0 – Release Date: May 2, 2018. All builds are cumulative. Applying 6.1.0 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 6.1

The 6.1 family of Turbonomic releases includes the following new features:

- User Interface Enhancements

We continue to work on how we present information about your environment. New things to see include:

- Dashboards Page – This release includes a number of additions to dashboards:

- * Executive Dashboards

Turbonomic ships with two executive dashboards in place – the On-Prem dashboard and the Cloud dashboard. The On-Prem executive dashboard shows the overall health, performance, and capacity in your on-prem datacenter. The Cloud executive dashboard shows your overall cloud expenditures and how you can improve performance and reduce cost.

- * Chart Widgets

This release includes a number of new chart widgets. In addition, the charts are more interactive. You can click legends to show/hide data, and individual charts include their own time-range selector. As you create dashboards, use the new Widget Gallery to choose and set up the charts you want.

- Color Scheme – We have changed the color scheme to make it easier to scan for and find the information you want.

- Scope and Groups – Now, when you use the Search page to set a scope, you can save the resulting scope as a group. You can use that saved group to scope policies or plans, and you can save it as a favorite scope for later searches.
- Supply Chain Navigator – For a given scope, you can use the supply chain to zoom in on specific tiers of entities in that scope.

- Enhancements to Action Recommendations

To improve the quality of resize recommendations for workloads on the public cloud, Turbonomic includes the current storage and network requirements. Even if the compute resources are underutilized on a workload, if the available templates cannot support the storage or network requirements then Turbonomic will not recommend a change.

For resizing VMs and Storage, this release introduces tuned scaling action settings. These settings give you increased control over the action mode for the affected actions. With this feature you can rely on automation for resize actions within a normal range (the tuned scaling range), and direct Turbonomic to post more conservative actions when the issue lies outside of the scaling range. For an overview of tuned scaling, see "Tuned Scaling Action Settings" in the User Guide.

In addition, we have clarified a number of the action descriptions that Turbonomic displays in actions lists.

- Public Cloud Enhancements

This release improves the display of costs on the public cloud, showing you:

- RI Inventory – See the count of RI workloads, account details, rate of utilization, effective cost, and savings for the RIs in your inventory.
- RI Utilization – The ratio of your RI inventory to the used RIs
- RI Coverage – The ratio of total workloads to RI matches in your environment

In addition, Turbonomic analyzes the history of resource utilization in your environment to recommend new RI purchases for the appropriate workloads.

- Plan Enhancements

This release adds two new plans to the planning wizard:

- Alleviate Pressure Plan – Use this plan to identify what workload to move from a cluster that is running "hot". It shows you the least amount of changes you should make to reduce pressure on the "hot" cluster.
- Optimize Cloud Plan – For cloud environments, see how to get the greatest savings possible while still assuring performance for all of your workloads. For example, this plan identifies opportunities to move workloads into RI purchase plans, as well as ways to reduce allocated resources without affecting application performance.

- Effective Host CPU

CPU cores and clock speed are not the only factors to determine processing power. For example, differences in processor architecture can result in CPUs with more or less effective capacity, given the same clock speeds.

For newer models of machines, the effective CPU capacity is a result of the number of cores, clock speed, and advances in processor design. Turbonomic includes a catalog of CPU models based on industry-standard benchmark data. This catalog maps CPU models to effective CPU capacity. To improve action recommendations, Turbonomic discovers the effective capacity for hosts in your environment and includes that information in its calculations. For planning, you can create host templates that include specific CPU models – The plan uses that information when calculating placement recommendations.

- Enhanced Container Support

This release includes the following enhancements to actions and automation for containers:

- Constraints – Turbonomic imports Resource Quota as constraints to refine scaling and placement decisions. Namespaces, Projects and Organizations, and Spaces appear as Virtual Datacenters in the supply chain. Cloud Foundry constraints include Memory Allocation and Number of Services for a Quota Plan. Kubernetes constraints include Memory Allocation and CPU Allocation for Quota objects.
- Container Resize – Turbonomic supports automatic and manual resize actions.
- Diego Cell Scaling – With Pivotal Operations Manager targets, Turbonomic supports Diego Cell scaling actions.
- VM Naming for Pivotal – With Pivotal Operations Manager targets, Turbonomic names VMs as follows: `{PCF Job name}#{index number}[{IaaS VM name}]`
- Container Plans – You can run plans to see how many worker nodes and cells you need to satisfy changes in Container or Pod demand.

- Enhanced Target Support

This version of Turbonomic introduces the following target enhancements:

- Cisco Tetration – Connect Turbonomic to the Cisco Tetration platform to discover network flow between workloads, and use keep critical network traffic on as short a path as possible.
- Cisco HyperFlex – This release introduces support for Cisco HyperFlex as a Hyperconverged target. HyperFlex targets add Disk Array entities to the supply chain, and receive more granular information from the compute resources in your environment.
- PaaS Targets – This release introduces support for Pivotal Operations Manager as a PaaS platform target. Integration with Pivotal Ops Manager adds Pivotal identifiers to the VM names, so you can use naming to identify these VMs and build groups. Taking advantage of Pivotal Operations Manager and BOSH, you can execute Diego Cell scaling actions. The Pivotal Ops Manager target also adds the related Cloud Foundry endpoint target, which enables discovery and management of the related resources.
- AppDynamics – This release introduces support for the AppDynamics platform. With an AppDynamics target, Turbonomic discovers the logical Business Application, and discovers the nodes that it contains.
- HPE 3Par – This release introduces support for Adaptive Optimization (AO). AO for HPE 3Par enables management of data storage across two or three tiers. AO places storage regions on the appropriate tier in response to periodic analysis that AO performs.

Update Recommendations

You can apply this update to any GA version of Turbonomic from version 5.9 or later, if it is running on CentOS.

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

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

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`

Storage Requirements for the Turbonomic Server

Turbonomic now states 150GB or greater as a requirement for disk storage. For Turbonomic servers hosted on VMware hypervisors, you should provide 150GB *plus* swap space to match the RAM allocation.

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.

Known Issues

- In public cloud environments, it is possible to create policies that enable workload moves across regions. Further, Turbonomic might recommend both a resize of the workload and a move from one region to another. If the resize is to a template that only exists in the new region, then Turbonomic displays an incorrect cost for the resize and move.
- After you run a Migrate to Public Cloud plan, you can then save the plan results to a PDF or a CSV file. For Migrate to Public Cloud, the Save As PDF does not work.
- For public cloud environments, under some circumstances Turbonomic displays incorrect values in the Cloud Cost Comparison chart.
To collect valid cloud cost statistics, the Hourly Saved Statistics must be set to a value greater than 24 hours. To check this setting, navigate to **Settings : Maintenance Options** in the user interface.
It is also possible that after you make changes to RI Cost, the values in this chart can be incorrect. If you make changes to **Settings : Budget and Costs : RI COST**, then you should rediscover your cloud targets, and wait 24 hours for the cost values to align with the new settings.
- When you set up a deployment to either reserve or deploy VMs, you can also specify workload placement policies to control the placement decisions. You specify workload placement policies for a given scope, where you give a group of consumers and a group of providers. However, the group of consumers must be a static group.
If you specify an empty static group of VMs as the consumers in a placement policy, and then use that policy for a reservation or placement, then the placement will not recognize the policy. To set up a policy for reservations and placement, the consumer group must be static, and it must contain at least one VM.
- In AWS environments, under certain circumstances Turbonomic can show negative costs in charts for specific regions. This can occur when the AWS target indicates that it offers templates, but those templates are not truly available. In that case Turbonomic calculates negative values for costs in that region.

- After you update a license, or you restart the tomcat server, you should manually rediscover all of your targets. This is to ensure that all the target data updates in your installation.
- The browser that you use for the Turbonomic user interface must be synchronized with the Turbonomic server to within one minute. Without this synchronization, Turbonomic can show incorrect metric values.
- When setting up a Migrate To Cloud plan, it is possible to choose a group of *cloud* VMs to migrate. You can run such a plan to simulate migration from one region to another. However, if you choose to migrate to a complete Cloud Service Provider (choose from Providers instead of Regions or Groups), then the plan will show inconsistent results. For example, the plan results can indicate that you have no VMs at all.
To plan for migrations across an entire service provider, you should use the Optimize Cloud plan.
- After you create a policy and set a scope to it, the user interface does not reliably take edits to the scope of the saved policy. If you need to change the policy, create a new one with the scope that you want.
- 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 a actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.
- If you use Google Chrome for the Turbonomic user interface, you must turn off the Chrome Preview mode before you download reports.
- 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.
- Under some circumstances in AWS environments, Turbonomic cannot recommend changing a workload from on-demand billing to an RI purchase plan. This occurs when the name, description, and platform fields of the workload's image do not clearly identify the OS. In that case, Turbonomic cannot identify the RI that matches the given workload.
To enable RI optimization for all of your AWS workloads, make sure the image fields include a clear OS name such as `ubuntu`, `centos`, or `Windows 10`. If your account can find an RI that matches the identified OS, then Turbonomic can recommend switching that workload to an RI purchase plan.

- Do not make policy settings on parent groups (groups of groups).
When you make policy settings on a parent group, Turbonomic resets the setting to the default every time you restart. This can have unintended consequences if you want these policies to implement business rules that control Turbonomic actions.
When setting up policies, be sure you set them to individual groups. If necessary, create a custom group for the settings you want to apply.
- In order to perform VM moves between Hyper-V servers in the same cluster, you must select `Discover Host Cluster` when adding one of the clustered hosts as a target.
- When using the Firefox browser, platform diagnostics cannot be downloaded. To download diagnostics, you must use another browser.
- The Deployer user type has access to all settings, instead of the subset of settings indicated in the User Management menu.
In Azure environments, storage cost depends on the type of storage you have purchased. For unattached, unmanaged standard storage, you are charged for utilized storage. However, Turbonomic calculates cost based on what storage has been allocated. As a result, calculated costs can be higher than actual costs.
In reports this can show in two ways:
 - Current costs can be overstated.
 - Potential savings for moves to other storage can be overstated.
- 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.

- In AWS environments, a Reserved Instance (RI) can be associated with a region or an availability zone. For RIs associated with zones, the Turbonomic scale actions assume reservations are regional. As a result, actions may resize a VM based on its reserved cost, but the VM charges will actually use the on-demand cost. If this occurs, use the AWS console to change the scope of the RI to "regional".
- 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 can show utilization on the remaining hosts as though the hosts to suspend are still active. The result is incorrectly low values for utilization on the hosts in the After Plan section of the chart.
- When working with templates in the JSON API, the field names for some of the input DTO fields do not match the corresponding field names in the template output DTO.
- 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`.
- Navigating to the Policy View will show default settings for entities not yet available in Turbonomic, and internally-named entity groups.
- When targets are added that introduce the Logical Pool entity to the Supply Chain, scoping to the Logical Pool entity will not show data for that entity.
- For Cloud Foundry environments, the Turbonomic market can incorrectly identify the VM that hosts a container. This can happen when two or more cloud targets have VMs with the same internal IP address. For a container that is hosted on a VM with that address, Turbonomic cannot assure that it hosts the container on the correct VM.
- 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.
- Under some circumstances in a NetApp environment, execution of a storage move action will fail.
- When you set up a Migrate Workload plan, you can choose the workloads to migrate from a list of VM GROUPS or a list of VM ENTITIES. When the user interface first shows the list of VM GROUPS, the list actually shows individual VMs, not VM groups. If you show VM ENTITIES and then show VM GROUPS again, the group list displays correctly.

- You should never use duplicate names for groups of the same entity type. However, the user interface does not validate group names to keep you from creating a duplicate name.
- To avoid timeout issues when working with saved plans, you should only delete one saved plan at a time.
- **Customer Issue 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.
- **Customer Issue 97962**

Automation policies for hosts include analysis settings for Desired State. The desired state settings affect Turbonomic analysis to give more emphasis to assuring performance or to increased utilization of infrastructure. You can specify Desired State globally in the Host default policy, or you can specify different Desired State settings for different scopes of hosts.

Turbonomic plans only use the global setting that you have in the Host default policy. The plans ignore scoped settings for Desired State.
- Turbonomic introduces a scaling constraint that you can use to exclude specific templates from move actions (VM, only) or resize actions (VMs and database instances on the cloud). For example, workload resizing in the cloud changes the template that defines the workload. To constrain resize actions for a scope of workloads, you can exclude certain templates for that scope.
- In Azure environments, when you set policies to exclude templates for resize actions on Database Instances, the template exclusion initially fails. If this occurs, you can force Turbonomic to recalculate actions, and those actions will then recognize your template settings. To do this, change the action mode for an action in this policy, and apply that change. Then reset the action mode to its original value and apply it again.
- 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 disabled for the Storage Defaults policy.

Find the Storage Defaults policy and make sure that the option for Disable Datastore Browsing is ON.
 - Make the following settings for the Virtual Machine Defaults policy:
 - * For **Action Automation**, set the following actions to **Manual**:
 - * Resize Up
 - * Resize Down
 - * Storage Move
 - * Start
 - * For **Resize Recommendation Incrementing Constants**, drag the Rate of Resize slider to **High**.
- 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.
- 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.
- To enable the best resize-down performance for VCPU on the cloud, you should set the global VCPU Resize Increment to 1 MHz.
- 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.

- To set up a Migrate to Cloud plan, you specify the scope of VMs that you want to migrate to the cloud. Currently, the plan configuration only supports a scope of one group or one VM. If you want to migrate workloads that are currently in different groups, then you can create a static or a dynamic group to include those workloads and migrate them to the cloud.
- Turbonomic supports logging in to AWS targets via AWS Identity and Access Management (IAM) with IAM Users or IAM Roles. To enable using Roles, you must run the Turbonomic software on an EC2 instance in the AWS cloud, and you must have the Turbonomic instance run as the IAM Role, and connect to the AWS target accounts with the appropriate IAM Role. To perform these actions, please contact Technical Support.
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.
- If multiple targets are added to Turbonomic with identical identifying fields, you will not be able to distinguish these targets from one another via the API or the user interface. To avoid this, ensure that you add targets with unique identifying fields.
- Customers upgrading from 5.7 or 5.8 versions of Turbonomic might need to reset their custom policies for disk arrays.
Starting with version 6.0, Turbonomic changes the way it creates groups categorized as Storage by Disk Array. In earlier versions, these groups included placeholder disk arrays (arrays that are not discovered through a disk array target). Turbonomic no longer includes these disk arrays in the Storage by Disk Array category of groups.
If you have upgraded from 5.7 or 5.8, and you used any Storage by Disk Array groups to set Turbonomic policy overrides to your disk arrays, then you will lose these settings as a result of this change to groups. You must reset these policy settings to your disk array groups in the upgraded version of Turbonomic.
- When you set up action orchestration in an Automation Policy, you should be sure that the scope is 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 Turbonomic can recommend different actions in real-time analysis and in plans. These differences can be for move actions in on-prem environments, and move or resize actions on the public cloud. If you see differences in these recommendations for plans vs. real-time analysis, check the VM Storage Move action modes and VM Host Move action modes for the affected scope. These action modes can be set to enable shared-nothing migration. Try disabling shared-nothing for on-prem, or enabling it for cloud. For more information about Shared-Nothing Migrations, see **Introducing Turbonomic: Turbonomic Actions : Shared-Nothing Migration Actions** in the User Guide.
- On rare occasions, the user interface for adding a target displays a blank page. If that occurs, refresh the browser to display the Add Target controls.

- 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.
- 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.

Fixed Issues

- **Customer Issue 100219**
Fixed: Under some circumstances in the public cloud, charts that display costs can include incorrect storage costs.
- **Customer Issue 100219**
Fixed: In Azure environments, under some circumstances the Azure target returns an empty list of regions, and Turbonomic displays an incomplete supply chain.
- **Fixed:** For Azure environments, because of changes to the Azure API, the pricing information that Turbonomic collects from Azure targets can be inconsistent.
- **Fixed:** The performance for the Hybrid and Cloud views of the Home Page is unacceptable.
- **Fixed:** A security scan against the Early Access version of Turbonomic returned unacceptable results.
- **Fixed:** In large environments, the performance when you navigate to the Search Page is unacceptable as the page loads the results for the **All** category.
- **Fixed:** In circumstances when any of the nightly plans fail, then subsequent discovery passes can also fail.

- **Customer Issue 100151**
Fixed: Under some circumstances Turbonomic can fail to discover a given HPE OneView target.
- **Fixed:** In environments with Virtual Datacenters (VDCs), if VDCs are removed then in some circumstances Turbonomic does not correctly discover the change. It displays an error that says `Cannot find replaced entity`.
- **Fixed:** When creating a group, the user interface can reject the group name as a duplicate name even if the name is just similar to another group name.
- **Fixed:** For integrations with vRealize Automation environments, under some circumstances a VM placement action does not discover all the available datastores, and the placement can fail.
- **Customer Issue 99986**
Fixed: If you have pending actions with a Manual action mode, and then create a policy for that scope to put the actions in Automated mode, the already pending actions do not change to Automatic.
- **Customer Issue 100030**
Fixed: In some circumstances after updating from earlier versions, Turbonomic displays a flickering user interface as you log in.
- **Fixed:** Users with the Administrator or Advisor role cannot delete saved plans.
- **Customer Issue 99908**
Fixed: Under some circumstances when discovering Tomcat servers, some metrics are not available and, as a result, the discovery fails.
- **Fixed:** Under some circumstances, Turbonomic VCPU and VMEM resize recommendations do not show complete data. For example, the "From" field is blank. In some circumstances, executing such an action fails to complete.
- **Fixed:** After using the Deploy Page to reserve workloads in your environment, the user interface shows the number of active VMs when you look at a given host, but the Details Page for that host shows the number of active VMs plus the number of reserved VMs. This is an inconsistency that can be confusing to the user.
- **Customer Issue 99824**
Fixed: The performance for the Pending Actions chart is unacceptable. It can also freeze the user interface.
- **Fixed:** For users with the Advisor role, when editing a plan the user cannot access the Change Automation Settings feature of the plan configuration.
- **Fixed:** For nightly plans, Turbonomic executes plans for ten clusters each night. In environments with more than ten clusters, plans and views rely on the generated data from the last run nightly plans. For example, for an appliance that had 21 clusters, the Exhaustion chart widget shows results from the last run nightly plans on that cluster. If the plan for that cluster has not been run yet, charts for some clusters can show a `No Data` message.
- **Fixed:** To create a plan, you can scope your session and then click the **PLAN** button. Then, as you configure the plan, the user interface should limit the lists of entities or groups to the scope that you started from, but it does not.
- **Fixed:** In the Classic User Interface, you can filter the list of actions to Show All or to Show Top. In some circumstances, Show Top does not include actions that should appear.
- **Fixed:** Charts for groups of groups (for example a group of clusters) do not show historical data from the database. They only show as much as two hours of data.
- **Customer Issue 99563**
Fixed: The Rightsize reports do not include any entries in their lists of pending actions.
- **Fixed:** Under some circumstances in AppDynamics environments, Turbonomic can post an excessive number of error messages to the log.
- **Fixed:** In AWS environments under some circumstances, Turbonomic recommends storage moves to avoid IOPS congestion when the IOPS utilization does not warrant that move.
- **Fixed:** The user interface now accepts edits to existing automation policies. For example, you can edit the name of an existing policy.

- **Fixed:** In some circumstances, a migration plan or a reservation can query a cloud target that has unavailable regions. In this case, the plan to calculate workload placement fails.
- **Customer Issue 96941**
Fixed: Under some circumstances, when browsing datastores for wasted files, Turbonomic can include active VSAN files in the list of wasted files.
- **Customer Issue 99264**
Fixed: Charts for magnetic storage (in public cloud environments) do not round values to two digits. As a result, the cost can display in scientific notation.
- **Fixed:** For AWS environments, the user interface displays an availability zone as a Host. An availability zone can have more than one storage resource, but the user interface only displays a single storage for such a host.
- **Fixed:** The JSON API enables you to create VM templates without specifying the storage type. If you use such a template with a plan, the plan will fail.
- **Fixed:** After restarting tomcat, some saved plans fail to reload in the user interface.
- **Fixed:** In the JSON API, you can pass the following call:

```
/markets/{market_uuid}/scenarios/{scenario_uuid}
```


The API succeeds if you pass the market name, but fails with a 400 if you pass a UUID for `market_uuid`.
- **Fixed:** When managing workloads in Kubernetes environments, Turbonomic uses Kubeturbo pods. If a Kubeturbo pod experiences a crash, then after redeploying it fails to connect with Turbonomic.
- **Fixed:** Under some circumstances, creating a Hardware Refresh plan fails with an error.
- **Fixed:** When setting VM policies in the user interface, the **Rate of Resize** setting can freeze on a given value.
- **Fixed:** After customizing the views on the Home Page and for standard dashboards, a Turbonomic update reverts to the default views.
- **Fixed:** The legacy user interface allows you to create an `administrator` user account that has scoped access. If you log in with such an account, then the Reports page does not show the list of standard reports.
- **Fixed:** Under some circumstances, Turbonomic generates excessive move actions.
- **Fixed:** The "Density and Headroom" chart widget has inconsistent display. In some cases, it shows a projection without historical data. In other cases, it shows historical data but no projection.
- **Customer Issue 99421**
Fixed: When you create a static group on the **Settings : Groups** page, you can use Search to filter the list of entities. Assume you filter the list by one string and select one or more entities. Then if you filter the list again, the user interface removes your current selection. As a result, you can't use successive searches to build up an accumulated selection of entities.
- **Fixed:** The JSON API does not support multiple tag values for a single tag category.
- **Fixed:** In the Reports Page, the list of generated reports does not scroll to list more items than can fit on the page's display.
- **Fixed:** Under some circumstances, API calls to get `notifications` fail.
- **Fixed:** In environments that use Raw Device Mapping (RDM), Turbonomic incorrectly recommends placement and resize actions for the RDM storage. Such actions are not valid for RDM storage.
- **Fixed:** After you execute actions that affect the market (a move action, for example) via the user interface or the API, subsequent calls via the JSON API to get `/rest/logs` will fail.
- **Fixed:** In AWS environments, under some circumstances Turbonomic can recommend moving from standard magnetic storage to a higher storage tier. However, AWS does not support automated moves of this type. For that reason, the Turbonomic move fails.
- **Fixed:** When setting up an Active Directory account, Turbonomic requires that you provide the Active Directory server, when a domain name should be sufficient.
- **Customer Issue 98685**
Fixed: The user interface for setting the dates in a schedule window displays a panel that covers some date

choices in the calendar, so you cannot set every legitimate date.

- **Fixed:** After running a Migrate To Cloud plan, the plan shows a VM count that includes inactive VMs, even though the plan does not migrate inactive VMs.
- **Fixed:** When you reload a saved plan, the supply chain shows topology from the real-time market and not from the saved plan market. For example, if the plan scope uses a dynamic group, and the group membership has changed since you saved the plan, then the supply chain will show the new group membership when you reload the plan.
- **Fixed:** When you edit a template definition, the user interface requires you to retype the template name.
- **Customer Issue 98065**
Fixed: Under some circumstances, Turbonomic recommends a move to address compliance issues, but, after executing the move, it still posts the compliance issue.
- **Fixed:** Under some circumstances, plans that have been run and saved fail to reload into the user interface.
- **Customer Issue 98164**
Fixed: Under some circumstances, Turbonomic does not save historical cluster information to the database, and the associated charts do not show complete data.
- **Fixed:** When creating schedule windows for automation, there are certain time ranges that the user interface does not allow. For example, if you set the time from 10:00 pm to 5:00 am the next day, Turbonomic posts an alert saying that the start time is greater than the end time.
- **Fixed:** When you reserve a VM on a host, when it's time to actually place the VM then Turbonomic uses the reserved resource allocation to calculate the placement. This can result in incorrect placement decisions. For example, if the reserved allocation is higher than the actual utilization, Turbonomic might decide the VM cannot be placed on the host.
- **Fixed:** The **Search** that you can use to filter different lists needs to be enhanced. If you type multiple terms, it interprets them as literal strings. For example, you can filter a list of actions by searching for `move`. But if you want to see all Move VM actions, you should search for `move vm`. If you enter that string, the list finds no matches, even if it includes VM Move actions.
- **Fixed:** Under some circumstances, Turbonomic posts a log message that warns of a circular relationship between host and storage. This warning is benign, and Turbonomic establishes the correct relationship. This message should not appear in the log.
- **Fixed:** In the new user interface, under some circumstances the chart display for data points does not match the tooltip for that data point. For example, the chart shows a value of 200, and the tooltip for that data point shows a value of 100.
- **Fixed:** Under some circumstances, the user interface can show an AWS Relational Database using resources from a Host entity, and Turbonomic can recommend actions to move Relational Databases to hosts. In the cloud environment, the database should use resources from a VM.
- **Customer Issue 97730**
Fixed: In the new user interface, the **Settings : User Management** page does not list existing user accounts.
- **Fixed:** If you deploy a Relational Database instance in AWS and then stop it, the Turbonomic user interface used to show that instance as active. Now, the Turbonomic user interface shows the instance as inactive.
- **Customer Issue 97734**
Fixed: Under some circumstances for Azure environments, Turbonomic fails to complete target discovery.
- **Fixed:** JSON API specifies units for Resource placements.
- **Customer Issue 97631**
Fixed: In some circumstances when specifying placement of templates, the storage type is not provided in the template. In this case, the plan to calculate the placement will fail.
- **Fixed:** The documentation now explains that for a Migrate Workload plan, the plan scope must include both the workloads you want to migrate, and also the providers that you want to migrate the workloads to. If you do not set the scope in this way, then the plan might not show any migration results.

- **Fixed:** When users refresh the browser, Turbonomic navigates back to the Home page. It should stay on the page the user was originally viewing.
Fixed: When performing discovery, Turbonomic checks the repository to make sure that the number of mapped entities and associated objects is reasonable for the given environment. However, this test also counts other data objects that are used internally for analysis. As a result, Turbonomic posts a warning that the repository has exceeded the allowed count of repository objects.
- **Customer Issue 97066**
Fixed: In some circumstances, Turbonomic analysis can show cyclical recommendations to provision more, and then less infrastructure capacity. This cycling can occur over a relatively short time frame; on a daily cycle, for example.
- **Fixed:** The JSON API returns links to other endpoints in the API, but these links use the HTTP protocol instead of HTTPS.
- **Fixed:** You should not create scheduled policies for one-time occurrence (RECURRENCE is set to Does Not Recur). If you create such a policy and save it, then you cannot edit the policy later.
- **Customer Issue 96965**
Fixed: In RedHat Enterprise Virtualization environments, under some circumstances Turbonomic recommends moving a VM across clusters, but the move action does not execute successfully.
- **Fixed:** In some circumstances for AWS environments that have Relational Database instances, Turbonomic starts up with errors related to the instance names.
- **Customer Issue 96375**
Fixed: When you enable a proxy via the `ipsetup` command, Turbonomic does not set a proxy for both HTTP and HTTPS protocols.

