Turbonomic 6.1.9 Release Notes

September 7, 2018

This document describes issues that are addressed in Turbonomic 6.1.9 – Release Date: September 7, 2018. All builds are cumulative. Applying 6.1.9 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 6.1.9 is available as an offline update. Please see the Green Circle article:

How To Perform an Operations Manager "Offline Update" - Latest Links Included

> **NOTE:** The OVA for this version of Turbonomic ships with an allocated disk storage capacity of 1 TB. The documentation for earlier versions states that the minimum storage requirements are 150GB.

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`

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:** Under some circumstances, Turbonomic does not roll up daily statistics into the database for historical data. This can result in reports that do not show all the requested data.
- **Customer Issue 101584**
  Fixed: In RHEV-M environments, under some circumstances if Turbonomic discovery fails to discover network links on a VM, then discovery for the RHEV-M environment fails to complete.
- **Fixed:** In Azure environments under some circumstances with multiple Azure subscriptions, Turbonomic discovers templates from one subscription but considers them viable for other subscriptions. As a result, some resize actions show unusually high investment costs.
- **Fixed:** In resource charts for a cluster, the chart shows a capacity for the full cluster, but shows utilization data for individual entities. As a result, the chart scale is too great for the charted data, and it is difficult to read the chart.
- **Customer Issue 101918**
  Fixed: Under some circumstances the user Reports Page does not display any generated reports, and users cannot access their generated reports. This can happen if there is a report filename that begins with a dot character.
- **Fixed:** In some circumstances Turbonomic cannot parse virtual HD names when it executes VStorage Move actions. As a result, it logs the actions as successful but does not succeed in executing the move.
- **Customer Issue 101969**
  Fixed: When managing applications via WMI, if you specify an application signature of `.*`, then Turbonomic displays excess utilization of resources for the affected applications.
- **Fixed:** In the user interface, many SEARCH fields include a **Filter** control. Hovering and clicking to display the filter is inconsistent and difficult to use.
- **Customer Issue 102066**
  Fixed: When setting up a plan to add workloads that come from clusters that are outside of the plan scope, the plan can fail to include all of the added workloads.
- **Customer Issue 102116**
  Fixed: In VMAX environments, if logical pools have very long names, then Turbonomic fails to discover the VMAX entities.
- **Customer Issue 102161**
  Fixed: In the user interface, charts fail to show data for the most recent two hours.
- **Fixed:** In VMAX3 environments, Turbonomic does not calculate values for Storage Amount Used.
- **Customer Issue 102288**
  Fixed: In Azure environments, when Turbonomic recommends a resize for a database instance, it does not properly calculate the required cores. The resulting recommended action can show a resize to zero cores.
- **Fixed:** Under some circumstances for large environments, discovery fails as Turbonomic assembles long lists of actions.

Known Issues

- When the Top Volumes chart is scoped to one Azure resource group, the chart also shows the volumes of another resource group.
- For AWS targets, the configuration includes a field for the Amazon Resource Name (ARN) Account ID. You should not provide a value for this field. If you provide an incorrect value for this field, then the user interface can show savings and investments for the wrong account.
- After an online upgrade completed, the User Interface continued to show the "Getting Ready" message. To remove the message and display the Login page, refresh the browser.
- If you configure a plan and then run it, then subsequent configuration changes will not take effect. For example, assume you configure a plan and run it. Then you make configuration changes and run the plan again. Not all of your configuration changes will take effect. Instead or changing and running the plan again, you should create a new plan with the configuration you want.
- Whenever you run a plan that calculates RI and other service costs in the public cloud, that plan uses the current settings you have for Enterprise Discounts and Discount Overrides. This is true even if you run a plan that you had created and saved in the past.
  For example, assume you have an Enterprise Discount of 50%, and then run a plan. The calculations will include the 50% discount. Now assume you remove the discount, and then run the same plan again. Now the plan will not include any discounts.
- For a Migrate to Cloud plan, when migrating from AWS to Azure the plan does not place Spot Instances from the AWS environment.
- 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.
- 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.
- 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.
- In Azure environments, Turbonomic can discover VM groups based on tags. However, if the tags use the square bracket character, then Turbonomic does not create the group. For example, Turbonomic does not create a group for the tag, `test:[MyTag`.
- By default, applications are Mission Critical. With this status, Turbonomic will not suspend the VMs that host them. You can set an application to a Normal priority so that, if there is a shortage of resources on the host, Turbonomic can suspend the VM that runs that application. If you change the application back to Mission Critical, either via scheduling windows or by resetting the policy, then Turbonomic will start up the application if it has been suspended.

- Azure provides certain regions that are not open for general public in terms of resource allocation and deployment. These regions are intended for specific Government entities. As of this writing, these regions include:
  - Australia Central
  - Australia Central2
  - France South

  Even though access to these regions is restricted, Turbonomic discovers these regions. However, Turbonomic does not include them in analysis or plans. For example, a Migrate to Cloud plan cannot migrate workloads to these regions.

- In public cloud environments, actions to delete storage volumes show an operational cost for the storage. This operational cost can be inconsistent.

- 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.

- For Azure environments, the Rate Card API for datacenters in Germany does not return metrics for OpenSUSE or RHEL operating system licenses. For Azure datacenters in Germany, Turbonomic does not display license costs for these operating systems, nor does it include those costs in its analysis.

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

- When you configure a plan that uses host templates to add hosts to the environment, you must ignore constraints. In the plan configuration, click **General** (in the plan configuration supply chain), and then turn on **Ignore Constraints**.

  The User Guide does not include this required step.

- **Customer Issue 100672**

  By default Apache no longer supports host names with underscore characters in the name. When you deploy Turbonomic, you should deploy it to a VM that does not include those characters in its name. If the host does include an underscore character, then Apache will respond with a 400 error when you try to open the user interface.

  If you cannot change the host name, then you can use a workaround in the Apache configuration file. For more information, contact Technical Support.

- 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.

- In AWS environments, under certain circumstances Turbonomic can show negative costs in charts for specific regions. This can occur when the AWS target indicates that it offers templates, but those templates are not truly available. In that case Turbonomic calculates negative values for costs in that region.

- In OpenStack environments, when you set up reservations to deploy workloads via OpenStack templates you must constrain the deployment to the OpenStack datacenter.

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

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

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

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

- Turbonomic policies include a default policy named **Global Actions Mode**. You can use this policy to globally disable all actions in Turbonomic. If you turn on the setting to disable all actions, then `Disabled` takes effect for all actions. However, in the default policies the action modes remain as you have set them. The user interface does not show that you have globally disabled these actions.

- For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.

- Turbonomic now supports deploying on M5 instance types through the AWS marketplace.

- 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.

  The user interface cannot display Suspend actions for storage unless you make the following action mode settings for storage:
  - `Delete` set to Manual
  - `Terminate` set to Recommend
  - `Suspend` set to Recommend

- 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 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.

- 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`.

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

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

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

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

- **Customer Issue 98320**
  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.

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

- In AWS environments, Turbonomic discovers data from a given AMI to determine the OS of an associated VM. The OS influences cost calculations.
  If you delete the AMI from your environment, then Turbonomic cannot discover the OS, and so it assumes a free Linux OS. This can result in incorrect calculations of cost in plans and in evaluations of real-time costs.

- After restarting the Turbonomic server, users must log into new sessions in order to continue using the user interface or the API.

- For Migrate to Cloud plans, when you migrate a VM that has an attached ISO image, the plan shows a move of a non-existent 0GB disk. This 0GB disk is a representation of the ISO image, and you can ignore the move action. The other plan actions for the VM are correct.

- For Tomcat, SQLServer, WebSphere, and other application or database targets that use a scope to identify target instances, Turbonomic can fail to validate or discover the targets. If you add a target via scope, and that scope does not have any VMs to host the target applications, then the target will not validate. If you later add hosts for the applications to that scope, Turbonomic does not dynamically recognize the change and then validate and discover. Even if you execute a Validate command for that target, Turbonomic can validate but it will not run discovery.
  To avoid this problem, make sure your applications are running on hosts before you configure the target. If you have encountered this problem (you added hosts to a scope after configuring the target), delete the target from Turbonomic and set a new target with this scope.

- When you set up action orchestration in an Automation Policy, you should be sure that the scope 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.

- 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.