



Turbonomic 6.4.16 Release Notes

April 3, 2020

This document describes issues that are addressed in Turbonomic 6.4.16 – Release Date: April 3, 2020. All builds are cumulative. Applying 6.4.16 onto any release of Turbonomic v5.9 or later will include all previous fixes. Please see the Turbonomic documentation for earlier versions of the Release Notes:

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

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

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

For Updates to the 6.4 Family of Turbonomic Releases

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

What's New for Version 6.4

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

- Scale IO Support

This release introduces support for Dell EMC ScaleIO software-defined storage. You can now configure targets to ScaleIO versions 2.x or 3.x.

Update Recommendations

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

NOTE:

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

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

Offline Update Links

Turbonomic 6.4.16 is available as an offline update. Please see the Green Circle article:

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

NOTE:

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

Related Green Circle Articles

- Offline and Online Update Instructions:
<https://greencircle.vmturbo.com/docs/DOC-1649>
- Release Notes and Product Documentation:
<https://greencircle.vmturbo.com/community/products/pages/documentation>

Configuration Requirements

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

Updating the Tomcat Server

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

After you update the Tomcat server:

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

Transport Layer Security Requirements

Starting with version 5.4, by default Turbonomic requires Transport Layer Security (TLS) version 1.2 to establish secure communications with targets. Most targets should have TLSv1.2 enabled. However, some targets might not have TLS enabled, or they might have enabled an earlier version. In that case, you will see handshake errors when Turbonomic tries to connect with the target service. When you go to the Target Configuration view, you will see a Validation Failed status for such targets.

In particular, we have found that NetApp filers often have TLS disabled by default, and that the latest version they support is TLSv1. If your NetApp target suddenly fails to validate after installing Turbonomic 5.4 or later, this is probably the cause.

If target validation fails because of TLS support, you might see validation errors with the following strings:

- No appropriate protocol
To correct this error, ensure that you have enabled the latest version of TLS that your target technology supports. If this does not resolve the issue, please contact Technical Support.
- Certificates does not conform to algorithm constraints
To correct this error, refer to the documentation for your target technology (for example, refer to NetApp documentation) for instructions to generate a certification key with a length of 1024 or greater on your target server. If this does not resolve the issue, please contact Turbonomic Technical Support.

Enabling HTTP and HTTPS Proxies

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

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

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

```
FLAGS="$ ( $JAVA_OPTS ) $CATALINA_OPTS "  
OPTIONS="-Dcatalina.base=$CATALINA_BASE ..."
```

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

```
-Dhttp.proxyHost  
-Dhttp.proxyPort  
-Dhttps.proxyHost  
-Dhttps.proxyPort  
-Dhttp.proxyUser  
-Dhttp.proxyPassword  
-Dhttps.proxyUser  
-Dhttps.propyPassword
```

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

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

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

For further assistance, contact Technical Support.

Enabling HTTP to HTTPS Redirects for the REST API

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

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

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

For more information, contact Technical Support.

Enabling IOPS and Network Monitoring for OpenStack Mitaka

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

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

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

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

Fixed Issues

- **Fixed:** For AppDynamics environments, under some circumstances Turbonomic fails to return lists of actions related to business applications, when the current specific user account is scoped.
- **Customer Issue 110226**
Fixed: For Hyper-V environments, under some circumstances discovery fails to discover the IPs of some VMs.
- **Customer Issue 109250**
Fixed: For vCenter Server environments, if you enable migration across vCenter instances, the migration should always keep the migrated VM on the same resource pool when applicable. For VMs managed by vCloud Director, migrations respect the resource pools. However, for environments that don't use vCloud Director the migrations don't always respect the resource pools.
- **Customer Issue 110125**
Fixed: In Hyper-V environments, under some circumstances when a resize down action fails for some reason, Turbonomic does not restart the VM (in its original size).

- **Customer Issue 110106**

Fixed: For Azure environments, when running a Migrate To Cloud plan, under some circumstances VMs with large disks fail to place in the "Consumption" plan for the migration. This occurs when the required storage capacity is only available on managed storage.
- **Customer Issue 110097**

Fixed: For DynaTrace environments, under some circumstances Turbonomic can discover incorrectly high heap capacity for a JVM. This can result in incorrect resize down actions.
- **Customer Issue 110093**

Fixed: Under some circumstances, logging unnecessarily shows the error:

```
CommodityStorageprovisionedVmStBoughtCommConverter
```

This can occur when a VM has RDM disks on several datastores. This is a normal situation, so the log should not post the error.
- **Customer Issue 110043**

Fixed: Cluster capacity calculations for headroom plans do not always match capacity calculations for the same cluster in other plans. As a result, plans to add workload might exceed cluster capacity sooner than predicted by the cluster headroom calculation.
- **Customer Issue 110063**

Fixed: For Horizon VDI environments, discovery can fail with the following exception: `IllegalArgumentException`. This can happen when the discovered entity contains special characters in its identifier. Discovery should not generally fail when it encounters a bad identifier, and discovery should handle special characters in the identifiers.
- **Customer Issue 110027**

Fixed: When you run a Migrate to Cloud plan, if the plan cannot place any of the migrating workloads, then the plan results show all the workloads that are currently placed in the cloud location. This is confusing.
- **Customer Issue 110011**

Fixed: Under some circumstances, the Health chart display applies the wrong set of issues to the given color. For example, assume 10 Without Risks, 20 with Minor Risks, and 5 with Critical Risks. The chart could show 20 green, 10 yellow, and 5 red.
- **Customer Issue 109843,109961**

Fixed: When creating groups based on filters, a given filter can find more entities than the group will include after you save it. This occurs because the filter is not case sensitive when you are creating the group, but the group itself uses case sensitive testing when assembling its members.
- **Customer Issue 109796**

Fixed: Turbonomic must improve heap management.
- **Customer Issue 102278,104021,107368**

Fixed: For Azure environments, under some circumstances an action to rescale a VM can fail because Azure does not return the necessary data on request. In this case, the VM can be left in a powered-off state. Turbonomic should return the VM to its original state.

Known Issues

- For very large environments that use the WMI targets, the WMI discovery can run out of memory. To address this issue, Turbonomic has improved memory handling with WMI discovery.

As stated in the *Target Configuration Guide*, Turbonomic recommends a maximum of 500 WMI entities per WMI target. If your WMI target manages more than 500 entities, then you can see further memory issues. If you must manage more than 500 entities per WMI target, please contact Technical Support.

- **Customer Issue 108374**

Under most circumstances, a scoped user cannot see actions for entities that are outside of the configured scope. However, when zooming in to Host entities, the user can see actions for storage that is outside of the user's scope.

- In Azure environments, you can create a VM that enables *Microsoft Azure Ultra Disk Storage* managed disks (UltraSSD). Turbonomic now discovers such VMs, and it will only scale such a VM to other instance types that also support UltraSSD. For each Azure subscription, Turbonomic creates a group of these VMs, so you can easily scope to the VMs that enable UltraSSD. Go to the Search Page, and search for groups with names that begin with `VMs_Ultra_Disk_Enabled_`. The name for each group begins with that prefix, followed by the name of the associated Azure subscription.

Under some circumstances, resize actions for VMs that enable UltraSSD can fail with the Azure error, `OverconstrainedZonalAllocationRequest`. As a result of this failure, the VM cannot restart after the attempted resize, and it remains powered down. We are currently working with Microsoft Azure Support to clarify this issue and understand the the circumstances that can cause it.

Preventing Resize Failures

To prevent this issue:

- Exercise caution if you use Turbonomic to execute the resize or move actions that it recommends for these VMs. If you do execute these actions, track the affected VMs to make sure they can restart. Be aware that you might need to recover from a restart failure.
- Create automation policies for these VMs that set resize and move actions to `Recommend` or `Manual`. After you upgrade to this version of Turbonomic, you can use the discovered groups to scope your policies.

Recovering VMs that Experienced This Failure

Earlier versions of Turbonomic did not distinguish VMs that support UltraSSD. As a result, it might have generated actions in the past to resize these VMs. If you executed such an action, or if the actions were automated for those VMs, then your VMs might be affected by this issue. To recover these VMs, you must reset them to their original size and instance type. To do this:

1. Identify the VMs with Failed resize or move actions.

Go to the Home Page for the Global Environment, and display the Hybrid view. Then scroll down to the **Accepted Actions** chart (be sure it is showing the Tabular view). Click **Show All** and look for VM Resize or Move actions that show a `Failed` state. In the action description, note the VM's name.

If the chart has a lot of entries, you can download the data and search in a spreadsheet. Or you can scope an **Accepted Actions** chart a discovered `VMs_Ultra` group.

2. For each VM that enables UltraSSD, and that has a failed resize or move action, inspect that VM in the Azure Portal.

Navigate to the Azure portal and search for the VM by its name.

3. Review the **Overview** information.

You should see that Azure failed to start the VM.

4. Find the original VM size and instance type.

In the Azure Portal, navigate to the **Activity log** and inspect the failed Start action. The resize shows the **Old Value** and the **New Value**. Note the **Old Value**.

5. Reset the VM size.

In the Azure Portal, navigate to **Size**. Reset the VM to the size that you noted in the step above.

- For Azure environments that include Ultra Disk Storage, by default VMs on this storage do not see move or resize actions. To enable these actions for the affected VMs, you must:
 - Enable Shared-Nothing Migration
 - Enable Shop-Together for multiple datastores

Set the action mode for VM Move and Storage Move to be the same; either Manual or Automated.

Turbonomic includes a configuration setting to enable this. For assistance, contact Technical Support.

To enable VM move and resize actions in your Azure Ultra Disk environment, please contact Technical Support.

- For public cloud environments, under some circumstances you can see conflicts between the costs shown in the Cloud Cost Comparison chart, and charts for individual cloud accounts.

RI Savings for Individual Accounts

Specific RIs can provide savings for multiple accounts. Charts for an individual account show the full RI savings, which can result in exaggerated savings for that account. If you add up savings in the affected individual charts, the result is greater than the overall RI savings you would see in the Cloud Cost Comparison chart. Note that setting the Cloud Cost Comparison chart to a global scope is the best way to view the total impact of Turbonomic actions on your environment.

Disabling Resize for VM Resources

To resize VMs on the cloud, Turbonomic uses Move actions to run the VM on a different instance type. For on-prem environments it uses resize actions for specific resources. If you want to disable resizing for public cloud VMs, you should disable move actions. If you disable specific resize actions for VMs on the public cloud, the resulting display of savings can be inconsistent between the Cloud Cost Comparison chart and other savings/investment charts for the affected scope.

Customer Issue 108966

In vCenter Server environments, when Turbonomic discovers a datastore, it then generates a group of the VMs on that datastore. If there are no VMs on that datastore, it does not generate a group.

If you later move a VM onto the given datastore, Turbonomic does not immediately create the group for that datastore. Instead, it waits for an event that causes a global calculation of groups. To force this action, you can call the API to execute a regroup.

- For AppDynamics environments, the user interface allows you to author a placement policy for Business Applications. However, you cannot save the placement policy. Turbonomic does not support placement policies for Business Applications.
- Turbonomic delivers with an option to display the Classic User Interface. This user interface was implemented in Adobe Flash. For various reasons, you might want to disable access to this user interface.

Turbonomic now includes the `vmt.legacy_ui_enabled` property in the `/etc/tomcat/tomcat.conf` file to show or hide the button that displays the Classic User Interface. To hide the Classic User Interface button, set this to `false`.

While the configuration setting removes the option to display the Classic User Interface, you can still navigate to this user interface by typing the URL directly into the browser. To ensure nobody can navigate to the Classic User Interface, delete the `/var/www/html/com.vmturbo.UI` directory from your Turbonomic installation.

If you have any questions, please contact your support representative.

NOTE:

When you update to a newer version, the update installs the Classic User Interface files, and resets the configuration setting. For each update, you will have to manually disable the display of the Classic User Interface, and delete the associated files.

- In the user interface, you can see entities of type Switch and type Fabric Interconnect. Note that Fabric Interconnect is a special type of switch. To create automation policies for these entities, set policies for the Switch entity type.
- For public cloud environments, the billing data that the public cloud service makes available to Turbonomic is not complete for the current day. As a result, the RI Coverage and RI Utilization charts can show incorrect values during this time. For example, the RI Coverage chart might show less than the actual coverage because it does not have the complete RI data.

This does not have any impact on the actions that Turbonomic generates.

To address this issue, set the time range to end on the previous day for your RI Coverage and RI Utilization charts.

- **Customer Issue 108198**

Fixed: For environments that result in large sets of actions, under rare circumstances Turbonomic can strand some actions in the queue of pending actions. This can occur when many new actions are added to the queue before the currently queued actions execute. To address this, you can modify your Turbonomic configuration. For more information, contact your support representative.

- The Monthly Savings or Investments Totals and the Cumulative Savings charts do not include RI Buy actions in their calculations.
- When editing automation policies, for Default Policies you can set the Rate of Resize for entity scaling. However, you cannot make a Rate of Resize setting for a scoped policy. The user interface includes a control for Rate of Resize in scoped policies, but any settings you make for a scoped policy will not take effect.
- In vCenter Server environments, charts can show that a Virtual Datacenter (VDC) uses resources at more than 100% of capacity.

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

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

The user interface displays this chart for real-time views, and it enables you to add this chart to dashboards.

However, this chart only shows real-time data for storage devices and storage groups. You should only rely on this chart when viewing Optimize Cloud plan results, or when your environment is specifically scoped to one or more storage devices.

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

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

Turbonomic discovers all the entities in the resource group except the locked storage. It also discovers the resource group.

- In Azure environments, when you run a Migrate to Cloud plan, the plan results do not include complete information for license costs for the migrated workloads.

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

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

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

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

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

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

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

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

IMPORTANT:

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

- When upgrading from a version earlier than 6.3.6, the upgrade can lose some policies that include user-created groups in their scopes. This is because the way Turbonomic manages group identities has changed. If you lose such policies, then you should create new groups and create new versions of those policies.
- When using the Classic user interface, under some circumstances the Inventory View does not update, or display content. If this occurs, perform a manual restart of the Turbonomic platform (restart the tomcat server), and the Classic user interface should perform normally.

- In AWS or Azure environments that include multiple billing families, if you scope the Turbonomic view to more than one billing family, the user interface does not display the billing family information correctly. For the **List of Billing Families** tab, it shows a count of all the accounts in the associated billing families. Also if you display that tab, the list can be empty, showing no billing families to inspect.

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

- **Customer Issue 105298**

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

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

- **Customer Issue 104680,107464**

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

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

- **Customer Issue 104812,104952**

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

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

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

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

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

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

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

- For Azure environments, the user interface can show storage entities for Standard SSD, even though these entities are not active in the environment. Turbonomic does not consider these entities in its analysis.

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

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

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

- For Azure environments that configure Consistent Resizing for a scope, if you run a Migrate To Cloud plan to migrate the VMs to AWS, then the plan actions do not maintain consistent resizing for the VMs.
- In AWS environments, for some templates the AWS API returns template values that do not match the AWS documentation. Specifically, the values for IO and Network bandwidth can differ from the values stated in the AWS documentation.

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

- After completing an online upgrade, the User Interface continues to show the "Getting Ready" message. To remove the message and display the Login page, refresh the browser.
- In AWS environments, availability zones can support different sets of instance types (templates). Turbonomic can identify valid templates for a given *region*, but not for a given availability zone.

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

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

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

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

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

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

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

- **Customer Issue 100341**

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

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

- **Customer Issue 100909,101615**

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

- **Customer Issue 100911,104082,104121,105565**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- In order to add an ACI Target, your Turbonomic instance must be using Market 1, which is disabled by default. If you require ACI Integration, please contact Turbonomic Technical Support.
- Turbonomic includes Automation Policies that you can create to modify the analysis of elements in your environment. This includes Scaling Constraints, which can include setting up templates to exclude when resizing workloads on the cloud. Note that the **Excluded Templates** option is not available for default policies for VM, Database, and Database Server templates.

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

- Customer Issue 99189,99805**

In vCenter environments, you might see unusually high storage latency values or excessive recommendations to provision new storage. There is a known problem with the storage latency values that vCenter Server versions 6.5.u1x and earlier return via the API. These versions can return unusually high storage latency values.

Turbonomic considers storage latency when calculating whether to move a VM to existing storage, or whether to provision new storage. Because of this known problem, Turbonomic can incorrectly recommend provisioning storage when moves are appropriate.

If you encounter this problem, then you should create a policy that disables storage moves for VMs that are managed by vCenter Server versions 6.5.u1x and earlier. To create this policy:

- Create a VM group that contains all the affected VMs. Note that Turbonomic automatically creates a group named `VMs_vCenter` that you might be able to use.
 - Create a new VM automation policy. This policy will disable storage move actions.
 - Set the group that you created to be the policy scope.
 - Under **Action Automation** add the `Storage Move` action and set it to `Disabled`.
- In cases where actions recommend that you suspend hosts, the Optimal Improvements chart should indicate no utilization on the hosts to be suspended. Under some circumstances, the chart can show utilization on these hosts. The result is incorrectly low values for utilization on the other hosts in the current scope.

- **Customer Issue 98910**

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

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

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

- Turbonomic generates special average or max utilization templates that it uses when calculating cluster headroom. You should not edit these templates, because Turbonomic will overwrite your changes the next time it generates the templates. However, the Template Catalog presents these templates as editable.
- After restarting the Turbonomic server, users must log into new sessions in order to continue using the user interface or the API.
- For Migrate to Cloud plans, when you migrate a VM that has an attached ISO image, the plan shows a move of a non-existent OGB disk. This OGB disk is a representation of the ISO image, and you can ignore the move action. The other plan actions for the VM are correct.
- For Tomcat, SQLServer, WebSphere, and other application or database targets that use a scope to identify target instances, Turbonomic can fail to validate or discover the targets. If you add a target via scope, and that scope does not have any VMs to host the target applications, then the target will not validate. If you later add hosts for the applications to that scope, Turbonomic does not dynamically recognize the change and then validate and discover. Even if you execute a Validate command for that target, Turbonomic can validate but it will not run discovery.

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

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

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

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

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

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

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