Turbonomic 6.2.2 Release Notes

October 12, 2018

This document describes issues that are addressed in Turbonomic 6.2.2 – Release Date: October 12, 2018. All builds are cumulative. Applying 6.2.2 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 6.2

The 6.2 family of Turbonomic releases includes the following new features:

■ Single Sign On (SSO) Support

You can now configure your Turbonomic installation to delegate authentication to a remote Identity Provider (IdP) that you have set up for your enterprise. This support uses SAML 2.0 to redirect login from the Turbonomic user interface to the IdP, and then to use the authentication tokens the IdP returns.

Turbonomic still applies the user roles you have configured to grant privileges to the user. For more information, see "Single Sign-On Authentication" in the *Installation Guide*, and see "Configuring a Group for SSO Authentication" in the *User Guide*.

■ Enhanced Container Support

For Kubernetes environments, this release adds support for configurations that declare node taints as a way to control Kubernetes scheduling, and the pods that use tolerations. Turbonomic recommends actions that respect these scheduling constraints.

■ Discovered Groups and Scope

Turbonomic discovers groups that you can use to set the scope of a session, policy, plan, or chart. This release introduces the following discovered groups:

- Azure Resource Groups

  In the context of Azure Resource Manager, a resource group is a container that holds related resources for an Azure solution. Turbonomic discovers the workloads that are in resource groups, and presents them as groups you can use to set scope. On the Search page, select **Resource Groups**. In the Optimize Cloud plan, you can set the plan scope to one or more resource groups.

- Accounts and Services

  Turbonomic now presents each AWS account and Azure subscription as a group that you can use to set scope. On the Search page, select **Accounts**. From the Top Accounts chart, if you click and account or subscription name, that also sets the session scope to that account. In the Optimize Cloud plan, you can set the scope to one or more accounts.

■ Public Cloud Enhancements

This release includes the following workload management improvements for the cloud:

- Improved Actions for Cloud Elasticity

  Analysis settings for VMs now include **Aggressiveness** and **Max Observation Period**. These settings help Turbonomic recommend more relevant actions, to better exploit the elasticity of the cloud. Instead of using the absolute peak value when considering VCPU or VMEM utilization, Turbonomic uses a percentile that you set as **Aggressiveness**. With **Max Observation Period**, you specify how much historical data to use when calculating the percentile.

  With these settings, Turbonomic evaluates the sustained resource utilization, and ignores bursts that occurred for a small portion of the samples. You can think of this as aggressiveness of resizing, where a lower percentile gives more aggressive resize actions. Recommendations better match the real-world utilization of your workloads, and the resulting actions should give you a more efficient, cost effective deployment. For more information, see "Aggressiveness" in the *User Guide.*

- Cloud Template Compatibility

  In Azure environments, the VM size for a template series includes a Maximum Data Disks (MDD) limit. For example, in the B-series the Standard_B1s MDD limit is 2, while for Standard_B2s the MDD limit is 4. Turbonomic recognizes these limits, and will not recommend moving an instance to a template that cannot support the current disk requirements. For example, Turbonomic will not move a Standard_B2s instance to a Standard_B1s template.

  In AWS environments, a VM can use Elastic Block Stores (EBS) or Instance Storage. If the VM's root storage is EBS, then Turbonomic will recommend a VM move. However, because Instance Storage is ephemeral and a move would loose

the stored data, Turbonomic does not recommend moving a VM that has Instance Storage as its root storage.

In AWS some templates require workloads to be configured in specific ways before they can move to those templates. If Turbonomic recommends moving a workload that is not suitably configured onto one of these templates, then it sets the action to Recommend Only, and describes the reason. Turbonomic will not automate the move, even if you have set the action mode for that scope to `Automated`. You can execute the move manually, after you have properly configured the instance.

The template requirements that Turbonomic recognizes are:

- Enhanced Network Adapters

- Linux AMI Virtualization Type

- 64-bit vs 32-bit

- NVMe Block

For more information, see "AWS Template Requirements" in the *User Guide*.

• Reserved Instance (RI) Purchase

For AWS environments that use *Instance Size Flexible* rules, Turbonomic can recommend that you buy multiple RIs of smaller instance types to cover specific workload needs.

Turbonomic tracks RI Coverage — The ratio of total workloads to RI utilization. By default, Turbonomic analysis drives toward 100% RI coverage. With this release, you can override this to set the coverage you desire. See "RI Purchase Profile" in the *User Guide*.

This release introduces a new Recommended RI Purchases chart that gives a breakdown of the RI templates that Turbonomic recommends for you to purchase. You can zoom in for details such as how the RI purchases map to workloads, cost, or where they will be placed. RI purchase recommendations include your total savings, and the estimated time to break even on the purchase investment.

• AWS Billing Families

Turbonomic now discovers AWS accounts that are consolidated into billing families. The Target Configuration user interface shows these families, including the master account and its various member accounts. For RI purchases, different accounts in a billing family can share the same RI resources. At the same time, accounts in other billing families cannot use those RIs. For more information, see "AWS Billing Family Recognition" in the *User Guide*.

• Cloud Discounts

Cloud service providers can offer discounts for workloads and services. With this release, you can specify a discount for an AWS billing family, or for an Azure subscription. You can also drill down into the discount to set discount overrides for specific services or template families. See "Cloud Discounts" in the *User Guide*.

- Cloud Planning Enhancements

  For this release we have improved the Cloud Cost Comparison table so you can more easily see the changes, savings, and investments the plan recommends for your environment. The plans also include the new Recommended RI Purchases chart. In addition, we introduced the following improvements for planning on the cloud:

  - Migrate to Cloud Plan

    On the cloud, VM templates usually include an OS platform to run processes on the VM. As you plan to migrate workloads to the cloud, you can specify the OS you prefer to run. You can keep the same OS that the original workload has, or you can choose to map the workload to a different OS. The plan can include the license costs for those OS platforms, or you can exclude the costs and treat the platforms as Bring Your Own License (BYOL). For more information, see "OS Migration Profile" in the *User Guide*.

    Migrate to Cloud plans support RI pricing in Azure environments.

    In accordance with Microsoft guidance on default storage, planning a migration to Azure environments defaults to using Azure Premium Storage for the migrated workloads.

  - Optimize Cloud Plan

    When setting the scope of the plan, you can now scope to:

    - Azure Resource Groups

    - Azure Subscriptions

    - AWS Accounts

    When planning for AWS environments, you can specify the RI Coverage that you desire. This is a way to fine tune the RI Buy recommendations that the plan generates.

    When you configure the plan, you can enable or disable RI Purchasing with the following settings:

    - Purchase RI and Optimize Workloads

    - Only Optimize Workloads

    - Only Purchase RI

- Online Documentation

  This release introduces an Online Help version of the User Guide and the Target Configuration Guide. To open the help, click the (?) icon at the bottom-left of the Navigation Bar, then choose Turbonomic Online Help.

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 6.2.2 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

■ **Customer Issue 102799**

Fixed: Under some circumstances in vCenter environments, the configuration of hosts can cause one aspect of discovery to fail, and that causes further discovery on that target to stop. As a result, Turbonomic does not discover all the entities in the affected target.

■ **Customer Issue 102816**

Fixed: When you set the scope to containers or container pods, the Multiple Resources chart does not show utilization data for the containers.

■ **Customer Issue 102561**

Fixed: In Azure environments, under some circumstances the Azure target can return null for a tagged database workload. In that case, Turbonomic incorrectly halts discovery for the affected target.

■ **Customer Issue 102144**

Fixed: In Azure environments, under some circumstances Turbonomic fails to discover all the metrics for entities in the environment.

■ **Fixed:** In Hyper-V environments you can designate VMs that have static memory capacity. However, for these VMs, Turbonomic recommends resize actions. For example, if a VM shows memory utilization approaching 100%, Turbonomic recommends resizing up even if the VM's memory is static.

■ **Fixed:** For Kubernetes environments, under some circumstances the Kubeturbo process can exit with an error.

■ **Customer Issue 102690**

Fixed: After running a Migrate To Cloud plan that migrates AWS workloads to Azure, the plan can show different costs in different results charts. This can happen when migrating AWS workloads that use RIs.

■ **Customer Issue 102639**

Fixed: When creating a dynamic group based on the memory capacity of the group members, the user interface for the Memory filter indicates that it uses MB. In fact, the filter uses KB. Must change the user interface to correctly show the units of measurement.

■ **Customer Issue 102648**

Fixed: In Dell Compellent environments, if you specify a custom port as part of the target address, Turbonomic fails to validate the target.

In earlier versions of Turbonomic, this used to work as the way to specify secure connections, or custom ports. In newer versions you choose an option to specify

a secure connection. But Turbonomic must still support custom ports in the target address.

■ **Customer Issue 102599,102920**

Fixed: In some environments with many different types of hypervisor targets, Turbonomic can execute VM moves excessively, appearing to cycle VM placement to different hosts.

■ **Customer Issue 101724,102157,102171**

Fixed: In some circumstances, Turbonomic recommends increasing memory for a database such that the database memory would exceed the hosting VM memory.

■ **Customer Issue 102584**

Fixed: In Hyper-V or VMM environments where multiple Hyper-V hosts use different datastores with the same names, Turbonomic can fail to discover the affected hosts.

■ **Customer Issue 102587**

Fixed: In vCenter environments you can configure Turbonomic to support VM moves across vCenter Server boundaries. To enable this, you create merge policies for the affected Host and Storage clusters.

However, you cannot create a policy that merges VSAN storage, because the VSAN does not appear in the list of Storage clusters. To fix this issue, Turbonomic must create a Storage cluster for VSAN Datastores that are not currently part of any Storage cluster.

The name of this cluster is the storage name, plus the word `Cluster`. (Note, this storage name is usually different from the display name in vCenter.) For example, if the storage is named `datastore-187` in vCenter, then the cluster is named `datastore-187Cluster` in Turbonomic.

■ **Customer Issue 102356**

Fixed: In vCloud Director environments, the vCloud Director target might be configured to manage multiple vCenter Server instances. Turbonomic must provide a way to specify a single vCenter Server instance for it to manage through the vCloud Director target.

■ **Customer Issue 102062**

Fixed: In environments that use placement policies to restrict one VM per host, if the VM includes multiple disks then move actions do not recognize the placement policy.

■ **Customer Issue 101479**

Fixed: In Turbonomic action scripts, you can execute calls before, during, or after actions execute to get information about the environment or to change action behavior. In these action scripts, the `$VMT_TARGET_UUID` variable does not correctly identify any targets.

- **Fixed:** When running Migrate to Cloud plans that migrate AWS workloads to another service provider (for example, to Azure), the migration plan does not migrate spot instances. However, the plan shows these spot instances as unplaced VMs.

- **Customer Issue 101330**

  Fixed: For NetApp 7-Mode environments, there are cases where the storage controller does not return a valid value for aggregated IOPS consumption. As a result, the following entities show IOPS of zero:

    - For Storage - Provided storage access = 0

    - For Disk Array - Consumed storage access = 0

    - For Disk Array - Provided storage access = 0

    - For Storage Controller - Provided storage access = 0

- **Customer Issue 102926,99478**

  Fixed: CPU Provisioned Used and Memory Provisioned Used settings made for a host do not take effect on the given host.

Known Issues

- In vCenter environments, to support VM moves across vCenter Server boundaries, the documentation instructs you to create merge policies to merge host clusters and to merge storage clusters. This enables Turbonomic to treat the merged clusters as one pool that it can move VMs into. In this way, you can move VMs from one datacenter into another.

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

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

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

    - First create the necessary merge policies for the affected datacenters

    - Then restart your Turbonomic

- When you configure a Cisco HyperFlex target, you can specify target credentials for the root user. For Turbonomic versions 6.2.0 and earlier, the root username must be set as `root`. For later versions, the username must be set as `local/root`.

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

- To enable Turbonomic access to AWS targets, you must log onto the target via an account that has specific permissions. The Target Configuration includes most of the required permissions, but it does not list the following two permissions that are necessary for read access to the target:
  - `organizations:ListAccounts`
  - `organizations:DescribeOrganization`

- For Dell Compellent targets, in earlier versions to specify a secure connection to the target, you provided the secure port as part of the target address. Starting with Turbonomic version 6.2.1 or 6.1.11, you specify a secure connection via the **Use Secure Connection** option in the target configuration.

  If you upgrade to the newer version of Turbonomic *and if you previously specified a port to set a secure connection*, then that target will fail to validate after you update. To validate the Dell Compellent target, edit the configuration to:
  - Remove the port specification from the **Address** field
  - Turn on the **Use Secure Connection** option

- 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 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.

- 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.

- **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.

- 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`.

- **Customer Issue 101794**

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

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

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

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

- 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.

- **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.

- 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.

■ 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.

■ 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 101735,99623**

  For AWS environments, while Turbonomic is discovering the environment, if it fails to reach one or more AWS regions, then all AWS discovery will fail for that target.

  Under some circumstance there are regions that Turbonomic cannot reach. For example, if you operate Turbonomic behind a firewall, then you might not be able to reach all the regions that are available to your AWS account. In that case, you need to specify which regions you want Turbonomic to discover.

  For information about how to specify the regions that you want Turbonomic to discover, contact your support representative.

- **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 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.

■ **Customer Issue 98910**

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

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

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

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

■ **Customer Issue 98320**

Customer Issue 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.