



Turbonomic 8.3.6 Release Notes

November 17, 2021

This document describes issues that are addressed in Turbonomic 8.3.6 – Release Date: November 17, 2021. Please see the Turbonomic 8 documentation for earlier versions of the Release Notes:

<https://docs.turbonomic.com/>

NOTE:

These release notes are for update to Turbonomic 8.3.6, from version 7.22.0 or later. For updates from the 7.21.x family or earlier, please contact your Technical Support representative.

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

<https://support-turbonomic.force.com/TurbonomicCustomerCommunity/s/customer-support>

What's New for Version 8.3.6

Version 8.3.6

There are no new features in this release.

Versioning Explanation

Turbonomic versioning uses V-R-M elements (Version, Release, Modification) in the version number to express the status of a given release, as follows:

Numbered Element	Example	Description
V - Version number	8.X.X	<ul style="list-style-type: none"> Changes to platform architecture or significant changes to data models
R - Release number	X.2.X	<ul style="list-style-type: none"> Major feature changes
M - Modification number is zero (0)	X.X.0	<ul style="list-style-type: none"> A quarterly release All preview features from previous bi-weekly releases are now GA No new preview features in this release
M - Modification number is greater than zero (1 or higher)	X.X.3	<ul style="list-style-type: none"> A bi-weekly release Can include new preview features Includes fixed issues

NOTE:

For API developers, the X.X.1 release can include final implementations of deprecated API features. These final implementations can make API changes that are not backward-compatible.

Configuration Requirements

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

Supported MariaDB Version

For its default historical database, Turbonomic currently supports MariaDB version 10.5.12. This support includes comprehensive testing and quality control for Turbonomic usage of the historical database.

If you are running Turbonomic installed as an OVA, and using the database that is included in that OVA installation, then you must use version 10.5.12. For versions of Turbonomic that you installed as an OVA *before* version 8.3.5, it is likely that you need to update your MariaDB.

For information about updating your MariaDB instance, see "Verifying your MariaDB Version" in the latest version of the *Turbonomic Installation Guide*.

Required DB Capacities

For Turbonomic versions 8.0.6 or later, your historical database must provide certain storage size capacities. For MariaDB or MySQL installations, you must ensure your database provides the necessary messaging and logging capacity.

If you installed Turbonomic as an OVA, and use the included MariaDB for the historic database, the easiest way to set the correct capacities is to update your Turbonomic to version 8.0.7 or later. For more information, see "Increasing your Database Capacities" in the latest version of the *Turbonomic Installation Guide*.

SQL Modes for External Databases

If you deploy Turbonomic to work with an external database instead of the included historical database, then you must specify the correct SQL modes for the database. Configure the database to support:

```
{ {ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION} }
```

In particular, the SQL modes should *not* include ONLY_FULL_GROUP_BY, NO_ZERO_IN_DATE, or NO_ZERO_DATE.

External DBs and Turbonomic Updates

If you have deployed Turbonomic with an external database server, for some updates you might need to manually create a new database and user for that deployment. This is important if your external database server is multi-tenant, or if your deployment does not grant administrative privileges to Turbonomic.

NOTE:

If your external database server is multi-tenant, or if your database server does not grant administrative privileges to Turbonomic, then you must continue with this configuration requirement.

Azure database services are multi-tenant. If you deployed an external database on Azure, this configuration requirement applies to you.

If you deployed your database server in a way that grants Turbonomic privileges to create new databases and new users, then the update will automatically create the required database. This configuration requirement does not apply to you and you do not need to take any action.

For some Turbonomic updates, the updated version includes new databases on the historical database server. If you are updating to one of these versions, then you must *first* create the new database, and a user account with privileges to access that database. After you have done this, you can update to the latest version of Turbonomic.

For more information, see "External DBs and Turbonomic Updates" in the latest version of the *Turbonomic Installation Guide*.

Transport Layer Security Requirements

By default, Turbonomic requires Transport Layer Security (TLS) version 1.2 to establish secure communications with targets. Most targets should have TLS 1.2 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 TLS 1.0. If your NetApp target fails to validate, this could be 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.

Improvements

- **Improvement:**

Savings and investments charts now reflect the effect of reverted actions.

Savings and investments charts can now stop accounting for realized savings or investments associated with a previously executed VM scale action as soon as that action is reverted via the cloud provider.

- **Improvement:**

Downloads for the Top Virtual Machines chart now include the full set of VMs.

When you download data from the Top Virtual Machines chart, the CSV file now includes all the VMs that are part of the chart's scope. In previous releases, the CSV file only includes a maximum of 50 VMs.

- **Improvement:**

This release improves performance for environments that can experience a flood of new entity creation.

Some environments can experience floods of new entities. For earlier releases, when this happened, the performance for discovery could become unacceptable. With this release, we have improved the performance for managing such floods of new entities.

- **Improvement:**

The Top Accounts chart now shows 30-day costs.

For cloud environments, the Top Accounts chart now shows the 30-day cost for each account. Cost covers compute, storage, IP, database, reservations, and data transfer services.

- **Improvement:**

For databases in Azure, the user interface now identifies when a database is geo-replicated.

For Azure environments, when Turbonomic discovers databases it can now identify ones that are geo-replicated. In this way, you can create groups that include only primary databases or replicated databases. For example, you can use those groups to create policies that treat replicated databases differently than primary databases.

Turbonomic uses the attribute **Replication Role** to identify the status of the database, where "Primary" is for a primary database, and "Secondary" is for a replicated database. You can use **Replication Role** in filters for search and Group.

Fixed Issues

- **Fixed Issue:**

For Database Server automation policies, you cannot add an orchestration workflow to the policy.

When you create an automation policy for a scope of Database Server entities, the user interface does not properly add Orchestration workflows to the policy. You can choose and add a workflow to the policy, but after you save

your changes the policy does not show that workflow in the user interface. In addition, if the policy is in effect, it is possible that you will get unexpected results.

If you edit an existing automation policy for Database Server entities, and that policy includes an orchestration workflow, when you save your changes the policy can lose the orchestration workflow, and it can give unexpected results.

- **Fixed Issue:**

Customer Issue 117655

For AWS, if you give an incorrect path for the AWS billing report, discovery can fail with an error.

For AWS environments, when you set up the S3 bucket to track metrics, you provide a path to the billing report. If you do not give a valid path prefix for the billing report, discovery can fail with an `array out of bounds` error. For example, this can happen if you give `root (/)` as the path prefix.

- **Fixed Issue:**

Customer Issue 117626

Analysis can recommend resizing a VMem reservation below the specified minimum.

Assume you set up a VMem resize policy that disables resizing below the VM's minimum. Under some circumstances analysis can recommend resizing below the VM's configured minimum. If resize down is automated, Turbonomic can take the action automatically.

- **Fixed Issue:**

For a persistent error condition, logging can flood with repeated messages.

For some errors that can occur when a probe discovers target metrics, the log can flood with repeated error messages while the error condition persists. The log should limit this flooding. For example, this can occur if connection problems result in a timeout or failure to get metrics from a target.

- **Fixed Issue:**

For VSAN storage, analysis can incorrectly treat Storage Access resize actions as executable, which results in an exception.

VSAN Storage Access resize actions should not be executable. However, analysis treats these actions as executable, and it then can result in an exception for the given action.

- **Fixed Issue:**

Customer Issue 117204

Charts for clusters can show incorrect utilization values.

Under some circumstances the charts for utilization on a cluster can show metric utilization for more than 100%. In this case, the cluster calculation does not correctly aggregate utilization across the cluster.

- **Fixed Issue:**

If discovery cannot access certain networks, deployment profiles can generate an error.

Assume an environment where Turbonomic does not have permission to discover all the networks related to a VM. In that case, the logs can show the error:

```
[DeploymentProfileMapper] : Scope entity: xxxEntityUuidxxx not found in EntityStore
for deployment profile
```

- **Fixed Issue:**

Customer Issue 117243

For Bare Metal environments, VCPU scaling actions can show an incorrect count of cores.

For Bare Metal environments, a VCPU scaling action can show an incorrect count of cores. The action does show the correct action results, and the correct MHZ for CPU.

- **Fixed Issue:**

For a VM that has multiple IP addresses, if there is an error resolving one address the log floods with the error message.

For a VM that has multiple IP addresses, if one address does not have a parent IP, then Turbonomic logs the error repeatedly. It should only log the error once.

- **Fixed Issue:**

For Azure, discovery posts an error for unsupported units for meter categories.

For Azure environments, discovery posts an error because it encounters unsupported units for meter categories in the price sheet.

- **Fixed Issue:**

Customer Issue 115844

For vCenter virtual datacenters (VDCs), the CPU Allocation and Memory Allocation charts show incorrect capacities.

For vCenter virtual datacenters (VDCs), the CPU Allocation and Memory Allocation charts show incorrect capacities.

- **Fixed Issue:**

When executing some actions, the platform can incorrectly post an error.

When executing move actions, under some circumstances Turbonomic can post an error similar to:

```
action-orchestrator 2021-08-04 16:02:50,167 ERROR [ActionDescriptionBuilder] :
getMoveActionDescription targetEntityId
Entity 73528895116576 doesn't exist in the entities snapshot
```

This is benign, and can result when analysis gets a list of actions that includes a move that has already been executed. The error should be changed to a warning.

Known Issues

- **Known Issue:**

For Kubernetes targets, the Target Status page shows when the KubeTurbo version does not match the Turbonomic version.

For Kubernetes environments, you configure the Kubernetes target via the KubeTurbo collector. The versions of the Turbonomic platform and the KubeTurbo agent should be the same. If the versions differ, the Target Status page shows a grey bar next to the affected target.

For more information about configuring KubeTurbo, and KubeTurbo notifications, see <https://github.com/turbonomic/kubeturbo/wiki/Troubleshooting#kubeturbo-health-notification-product-version>.

- **Known Issue:**

Customer Issue 117307

When updating your installation, under rare circumstances the Topology Processor component can fail to restart.

When updating Turbonomic, under rare circumstances the `topology-processor` pod can fail to restart. The log posts an error with the statement:

```
AccessDeniedException: /home/turbonomic/data/kv
```

If you experience this problem, please contact your support representative for a workaround.

- **Known Issue:**

For Azure, targets in the Australia Central region can show inconsistent pricing for VM license costs.

For Azure environments running in the Australia Central region, the pricing for license costs that is reported to Turbonomic can be incorrect. As a result, the user interface can display incorrect license costs for Linux or Windows operating system licenses.

- **Known Issue:**

For Azure and AWS, analysis does not consider the full cost for some workloads.

For Azure environments, analysis considers that base OS cost, but does not consider additional costs for support or other add-on features that are bundled with the OS. The affected OS types are Ubuntu PRO, SUSE 24/7, and RHEL with HA.

For AWS environments, analysis does not consider AWS Marketplace costs.

- **Known Issue:**

Customer Issue

For Fabric environments, when you scope the view to Datacenter the supply chain can fail to include all the related Host entities.

For environments that include Fabric targets, under some circumstances a view that you scope to the Datacenter might not include all the related Host entities.

- **Known Issue:**

Customer Issue 117106,117174,117545

After updating to a new version, you can experience poor performance when restarting the platform.

After you update Turbonomic, you restart the platform. If you are updating from a version that is earlier than 8.2.3, the restart can take an unusually long time. For such an update, the platform performs internal calculations to enable certain improvements that were introduced in version 8.2.3. While the restart is in progress, the user interface will display incomplete data.

If this problem persists, or Turbonomic shows ongoing problems during restart, contact your support representative for resolutions to this problem.

- **Known Issue:**

For very large container environments, when you update the platform you can experience excessive downtime for the historical data component.

When you update Turbonomic from version 8.3.2 or earlier to version 8.3.3 or later, if you have a very large container environment you can experience excessive downtime (up to three hours) for the history component. This can occur as the database performs a migration that improves overall product performance.

To avoid this downtime, you can run a SQL script to perform the database migration before you perform the product update.

1. Save a copy of the script to your Turbonomic installation.

- a. Log in to the Turbonomic VM.

Use SSH to log in to the Turbonomic VM using the turbo account and password.

- b. Change to the scripts directory.

```
cd /opt/local/bin
```

- c. Get the latest version of the script.

```
curl -O --proxy PROXY_NAME_IP:PORT \  
https://\  
raw.githubusercontent.com/turbonomic/t8c-install/master/bin/833update-vmtdb-  
cnt-table-index.sql
```

Where `--proxy PROXY_NAME_IP:PORT` is an optional specification to execute the download through a proxy.

- d. Save the script on the Turbonomic server to:

```
/opt/local/bin/833update-vmtdb-cnt-table-index.sql.
```

2. Run the script to execute the database migration.

Note that the database migration can take up to three hours to complete.

While still in the SSH session, execute:

```
mysql --user="username" --password="pwd" --database="vmtdb" < "./833update-vmt  
db-cnt-table-index.sql"
```

Where `username` is the turbonomic administrator username, and `pwd` is the password for that account. Note that `vmtdb` is the name of the database you will migrate.

For assistance or more information, contact your support representative.

- **Known Issue:**

For certain installation configurations, an online update can take longer than expected.

If you have installed Turbonomic behind a firewall and opened ports for the Docker Hub services that deliver the Turbonomic, it is possible that the script `onlineUpgrade.sh` can take 30 to 45 minutes to complete. If you have trouble running the update, please contact your support representative.

- **Known Issue:**

If your installation uses a Master Key Kubernetes Secret, you should back up the key data.

For versions of Turbonomic that you installed as version 8.1.5 or later, you should be sure to back up the Master Key data that is stored in your platform as a Kubernetes Secret.

The installation procedure creates a Master Key secret in the Kubernetes cluster. Turbonomic uses this secret to provide access for the platform components. You should save the key data to a safe location. If for some reason the key data gets corrupted or is otherwise unusable, Turbonomic will fail to operate. If this happens, you can contact your support representative and use this saved data to recover your platform.

To save the data:

1. List the platform secrets.

Execute the command:

```
kubectl get secrets
```

The results should include the Master Key secret, similar to the following:

```
...
  master-key-secret          Opaque          1          57d
  ...
```

If you see an entry for `master-key-secret`, then you should display the data and save it.

2. Display the Master Key data.

Once you find the Master Key name, you can then display the key data:

```
get secret master-key-secret -o yaml
```

The command result should be similar to the following:

```
apiVersion: v1
data:
  primary_key_256.out: AfnJWutxNHAduaIOdAii3DRA2fMa6lzX4rWetZxxZvc=
  ...
```

The important key data to save is the `primary_key...` data. In the above example, you should save the line:

```
primary_key_256.out: AfnJWutxNHAduaIOdAii3DRA2fMa6lzX4rWetZxxZvc=
```

3. Save the data to a safe place.

Write this data to a file and save it in a safe backup location. If you ever need to recover the Master Key, your support representative will use this data to perform the recovery.

- **Known Issue:**

After updating the platform, Embedded Reports can fail to display.

Under some circumstances, after you update Turbonomic to a new version, the Embedded Reports page can fail to display. The update appears to have completed successfully, but the Embedded Reports components appear to be running and ready. However, when you click on the Embedded Reports button, you get the following error:

```
failed to log in as user, specified in auth proxy header.
```

This can occur when the host VM takes a long time to completely start up all the platform components. If you encounter this problem, make sure all the components are running, and then execute the following command to restart the grafana pod:

```
kubectl delete pod -l app=grafana
```

For assistance, contact your support representative.

- **Known Issue:**

Updates from version 7.22.6 to version 8.2.0 can fail to complete.

If you are updating Turbonomic from version 7.22.6 to version 8.2.0, *and* you have enabled Embedded Reporting, then the update can fail to complete. If this occurs, recover the snapshot of your original installation and try to update to an earlier version (for example, 8.1.6). Then try to update that version to 8.2.0.

If the problem persists, please contact your support representative.

- **Known Issue:**

Customer Issue 116930**For Fabric environments, when you scope the view to Datacenter the supply chain can fail to include all the related Host entities.**

For environments that include Fabric targets, under some circumstances a view that you scope to the Datacenter might not include all the related Host entities. This can happen for hosts that include the hyphen character ("-") in the Host name. For example, for Cisco UCS targets, the supply chain will not show UCS hosts (that use the hyphen character in their names) when you scope to the Datacenter entity.

- **Known Issue:**

Customer Issue 115635**For Azure, under some circumstances a successful scaling action appears in the logs as failed.**

For Azure environments that include Availability Sets, under some circumstances a successful scaling action appears in the Turbonomic logs as failed.

- **Known Issue:**

Customer Issue 114942**With Executed Actions charts, some data is missing for actions on entities that have been removed from the environment.**

When you view Executed Actions charts or export data from them, some data is missing for actions on entities that have been removed from the environment. For example, assume an action was executed on a storage volume, and

that volume has later been removed from the environment. In that case, the exported data for that action will not include values that describe the removed volume.

- **Known Issue:**

The Onboarding wizards can sometimes fail to close.

When you first install Turbonomic, the user interface displays onboarding wizards to walk you through setting up your license, and configuring your first target. Under some circumstances, the button to end the wizard's workflow does not close the wizard. That can result in blocking you from continuing your Turbonomic session.

If the onboarding wizard does not close when you click **End Setup**, refresh the browser. That should close the wizard and leave you on the last user interface page that you visited.

- **Known Issue:**

For policies that exclude certain cloud tiers, when the cloud provider adds new tiers, they can appear as included in the policy.

For public cloud environments, when you make a policy to include only certain tiers (VM or Storage types) of entities, if the service provider deploys new tiers, then those will also be included in your policy.

This can be unexpected. For example, assume you create a policy to include only one VM type. Then if your service provider introduces new VM types, your policy will subsequently include those new types.

You should periodically check your policies to see if new tiers have been added to the INCLUDE list.

- **Known Issue:**

For Kubernetes environments, when you enable Feedback and Diagnostics for your installation, the collected data can include Kubernetes cluster names.

To help us improve the product, you can enable Turbonomic to collect anonymized and non-confidential data as you go about using the product. However, because of the way Kubernetes discovery works for Turbonomic, the collected data includes the names of any Kubernetes clusters that you have set up as targets. We do not make use of those cluster names in any way.

If you do not want Turbonomic to collect these cluster names, then you can navigate to **Settings / Maintenance Options / Feedback and Diagnostics** and turn off the option to share anonymized usage data.

- **Known Issue:**

Changes to a policy do not immediately show up in the user interface view of the affected scope.

When you set the scope of the Turbonomic view to a group, you can then view the automation policies that impact the given group. If you edit a policy for that group (in Settings: Policies), and then scope the view to that group again, the policy changes do not appear in the display for that group.

The display should update within ten minutes, after the next round of incremental discovery. If the condition persists, log out of your session and log in again to update the display.

- **Known Issue:**

Customer Issue 113340

Hitachi Vantara targets can incorrectly show multiple compliance actions.

For Hitachi Vantara environments that enable storage replication, Turbonomic can generate repeated actions of the type, `Reconfigure Storage to Provide Extent`. In addition, replica storage entities and their providers can show incorrect utilization values.

- **Known Issue:**

When creating an automation policy, you can assign two or more schedules to the Action Execution Schedule setting. If the action is one that cannot be executed by Turbonomic, then for Service Now integrations the action appears in the Change Requests for display, only. Also, such an action is *only* affected by the first schedule in the policy.

- **Known Issue:**

For Migrate to Cloud plans, under rare circumstances the plan's actions list can show duplicate entries.

For Migrate to Cloud plans, under rare circumstances the plan's actions list can show duplicate entries.

- **Known Issue:**

For Azure environments, discovery does not support the Brazil Southeast region.

For Azure environments, Turbonomic does not discover the Brazil Southeast region. Azure provides this region only to give business continuity and disaster recovery to workloads in Brazil South that require data residence.

The user interface does not display the Brazil Southeast region in any lists or charts. Also, if you do have workloads on that region, Turbonomic will not discover those workloads.

- **Known Issue:**

For ServiceNow environments, Turbonomic fails to save any automation policy that sets the Action Type to **Request Approval from ServiceNow**.

- **Known Issue:**

For AppDynamics environments, the platform cannot discover databases if the target authentication uses oAuth for credentials.

For AppDynamics environments, Turbonomic cannot discover databases if the target authentication uses oAuth for credentials.

- **Known Issue:**

For Application Component automation policies, the user interface allows you to make conflicting settings.

The Action Generation setting can show incorrect values that you can choose for the policy. As a result, you cannot save the policy.

- **Known Issue:**

For ServiceNow environments, the user interface allows you to set orchestration for actions that the ServiceNow integration does not support.

For ServiceNow environments, the Turbonomic user interface allows you to set orchestration for actions that the ServiceNow integration does not support. If you configure orchestration for these actions, then either the action never generates a ServiceNow CR, or the action can fail when the CR is approved.

The actions you can configure but will not generate a CR are:

- Storage Suspend
- VSan Storage Resize

Note that storage resize for a VSan is accomplished by provision/suspend of Host. You should not configure ServiceNow orchestration for VSan Storage Resize. However, Host Provision is not currently supported for ServiceNow orchestration (see next).

- Host Provision
- File Delete
- Application Component - No actions are supported

The actions you can configure but that can fail include actions that you must also configure for execution on the affected targets. These actions include:

- Host Suspend

For this action to succeed, it must be enabled in the given hypervisor, and there must be no VMs currently running on that host.

- Storage Provision

Currently Turbonomic can only execute a CR for this action on Pure and Dell Compellent storage.

- **Known Issue:**

The user interface does not currently show the billed costs for some Azure resource groups.

For Azure environments, when you inspect resource groups, Turbonomic does not currently show the billed costs for those resource groups.

- **Known Issue:**

Customer Issue 111396

For cloud environments, under rare circumstances analysis can recommend resizing a VM to an instance type that is older and less capable than an equally priced instance type.

Under most circumstances, when a cloud provider offers a new instance type that is meant to replace an older type, the provider offers it at a lower cost. In at least one instance we have seen a case with identical costs for the newer and older instance types. If this occurs, and capacity and cost are equal, Turbonomic cannot ensure that it chooses the newer instance type.

To work around this issue, you can create an Action Automation policy that excludes the older instance type.

- **Known Issue:**

The All Actions chart does not include pending actions for databases or database servers.

The All Actions chart does not include pending actions for databases or database servers.

- **Known Issue:**

Customer Issue 110123

There is a memory limit for the data you can download from the All Actions chart.

There is a memory limit for the data you can download from the All Actions chart. For example, assume you have executed many actions over time in your environment. As a result, the list of all executed actions might exceed the data limit. In that case, downloading a CSV file from the All Actions chart will fail.

- **Known Issue:**

Under rare circumstances, the etcd.service can fail.

Under rare circumstances the Turbonomic platform stops responding. This occurs when `etcd.service` fails. When it does occur, you should see the following error:

```
Error response from daemon: endpoint with name etcd1 already exists in network host
```

To recover from this situation, restart the docker service for the Turbonomic platform. execute the command: `sudo systemctl restart docker.service`

- **Known Issue:**

You must use certain templates when using PLACE to set up reservations or deployments.

When you use the **PLACE** page to set up a reservation or a deployment, you choose the templates to represent the workload you will deploy. The templates you choose must include an **Image** specification that gives the path to the VM package, and optional placement constraints.

Typically, you will use templates that are discovered through your hypervisor targets. Along with discovering resource capacities for the given VM, Turbonomic should also discover the Image specification for a given discovered template. However, in this version Turbonomic does not discover the Image descriptions. In addition, discovered templates and their image specifications are read-only. For this reason, you cannot set up placement or reservations using discovered templates.

- **Known Issue:**

For resources that do not support Reserved Capacity, charts can show them with zero reserved capacity.

Ring charts that show the utilization of different resources show a yellow segment whenever the Reserved Capacity for the resource is zero. For some resources there is no concept of reserved capacity, yet the ring chart still shows a yellow segment.

- **Known Issue:**

Optimized Improvements for plans do not include hosts to provision.

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

- **Known Issue:**

Customer Issue 99189,99805

In vCenter environments, you might see high storage latency, or excessive storage provision.

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`.
- **Known Issue:**

The Optimal Improvements chart can show incorrect data for hosts to be suspended.

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.

- **Known Issue:**

For vSAN environments, under some circumstances a plan that is scoped to a datacenter can fail.

For vSAN environments, when running plans that add or replace hosts to the environment, under some circumstances the plan can show the incorrect count for hosts, and the plan can fail.

This can happen for plans that meet the following conditions:

- The plan type is Hardware Refresh, Add Workload, or Custom
- The plan scope is set to a datacenter, and it includes vSAN hosts
- The plan uses an HCI template to replace the hosts

After running, the plan shows the full count of hosts in the vSAN environment, instead of the count of hosts in the plan scope.

To avoid this situation, do not scope the plan to the datacenter.

- **Known Issue:**

Customer Issue 105693

The Headroom chart for All On-prem Hosts does not always agree with the Top Clusters chart.

The Headroom chart for All On-prem Hosts does not always agree with the Top Clusters chart.

Turbonomic generates the All On-prem Hosts headroom data in a nightly plan. When the plan runs, this data is correct. In the course of the day, this data can become stale.

To accurately track your cluster usage, you should use the Top Clusters chart.

- **Known Issue:**

Customer Issue 107699,108212,109349,110474,110896,112196,112639,112749,112790,113104,114386,115748

Hardware Refresh plans to replace hosts can give unexpected results.

When you run a plan to replace hosts in a cluster, the results can incorrectly show that you need more hosts than you would expect. This can occur for two reasons:

- If the host template for the replacements does not use CPU specifications from the catalog, then the calculations for host capacity can be incorrect.

- When placing VMs on the replacement hosts, analysis assumes all the VM peaks can occur at the same time. This results in excessive utilization of peaks capacity in the plan.

To configure a plan in a way that avoids these problems, you can:

- Make sure the plan scope is for a single cluster.

This is the typical use case for a Replace Hosts plan. Record the cluster name, because the name will come in handy in later steps.

- Always use the **Select from Catalog** option when you create your Host template.

If the catalog does not include the CPU specifications that you want, choose an entry that is as close as possible.

To adjust the number of cores, you can then specify a different number of sockets. For example, assume you choose a CPU spec that includes 4 cores of a given core frequency, but you actually want 32 cores. You can choose that spec, and then set **Sockets** to 8 to achieve that number.

- Turn off **Scale** for the VMs.

After you choose the host template to use, click **NEXT: VIRTUAL MACHINE ACTIONS**. Then turn off the **Scale** option.

When replacing hosts, it's important to *not* scale the VMs, so you can see how the hosts can support your given workload.

- Replace all your VMs with a Cluster Average template.

Click **SKIP TO CONFIGURATION** to display the plan configuration, and open **Replace/Virtual Machine**. Display **Clusters** and click **Select all**. Then click **NEXT** to choose the VM template.

To choose the VM template you will replace with, type your cluster name in the Search box. The Templates list will show an AVG template for that cluster. For example, if your cluster name is MyCluster1, then the template name would be `myDomain.com: :AVG:MyCluster1`. This template captures the average VM utilization over the last ten days.

Choose this template and click **SUBMIT**.

- You can now run the plan.

- **Known Issue:**

For vCenter Server environments, Turbonomic does not recognize DRS rules for VM restart dependencies that are based on ClusterDependencyRule.

For vCenter Server environments, Turbonomic does not recognize DRS rules for VM restart dependencies that are based on `ClusterDependencyRule`.

You might be able to achieve a similar effect by expressing dependencies via `ClusterVmHostRule` or cluster affinity or antiaffinity rules.