



Turbonomic 7.22.5 Release Notes

July 27, 2020

This document describes issues that are addressed in Turbonomic 7.22.5 – Release Date: July 27, 2020. Please see the Turbonomic 7 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 7.22.5

Version 7.22.5

- **Action Visibility for Applications**

To assure application performance, Turbonomic recommends actions on the underlying *Application Components* and *nodes* (such as database servers, VMs, and containers) that provide resources to your Business Applications. The Pending Actions chart at every level of the application model shows these actions, thus providing visibility into the underlying risks that have a direct impact on your Business Applications' performance.

NOTE:

You can set automation and constraint policies for these underlying entities as usual.

- **Improved Workflow for Creating Automation Policies**

This release introduces a more intuitive workflow for creating automation policies. It is now easier to set action *automation modes*, *schedules*, and *orchestrators* (such as ServiceNow).

You can now assign a schedule to action execution, as a way to defer execution to a non-critical time window. For example, if a workload experiences memory bottlenecks during the week, you can defer the necessary resize to the weekend. Even if the workload has minimal utilization over the weekend, Turbonomic can recognize the need to resize, and will execute the action.

For details, see "Creating Scoped Automation Policies" in the *User Guide*.

- **Native Integration with ServiceNow**

This release introduces *Turbonomic Actions* for versions 7.22.5 and beyond. This is a ServiceNow application that you install via an update set. Turbonomic Actions integrates Turbonomic with your ServiceNow account to log the actions Turbonomic generates, and to defer these actions to your ServiceNow workflow for approval. Upon approval, Turbonomic executes the actions at the scheduled time, and logs the result in ServiceNow.

Turbonomic Actions uses the **Execution Schedule** that you assign to actions in Turbonomic automation policies. For information about setting up policies for these actions, see "Creating Scoped Automation Policies" in the *User Guide*.

To get the Turbonomic Actions application, contact your sales representative.

- **Horizon VDI Used Sessions**

Turbonomic now tracks two states of used sessions for your Horizon VDI desktop pools - *Active* and *Disconnected* (non-terminated). To view relevant data, scope to a Desktop Pool or View Pod and see the following charts:

- (*New) Total Sessions chart – Plots used sessions (Active and Disconnected) over time
- Active Sessions chart – Plots Active sessions over time
- Capacity and Usage chart – Shows current Total and Active sessions

This release also introduces a new *Desktop Pool Active Sessions* search filter for VMs. Setting the value to TRUE returns a list of VMs on your desktop pools with active sessions, while a value of FALSE returns all VMs with used sessions.

- **Entity Placement Constraints**

When you drill down to a single entity, you can see details about the entity's relationships in the supply chain. This shows you which entities provide resources to the entity, and which other providers can be available to it. For each listed entity, you can open a **Constraints** fly-out that gives more details about limits on the current entity's supply chain relationships.

For details, see "Entity Placement Constraints" in the *User Guide*.

- **Embedded Reports**

This release introduces Turbonomic Embedded Reports. Use this feature to understand application resource management trends, and to share insights with stakeholders via reports and dashboards.

Embedded reporting runs in its own component, as part of the Turbonomic platform. This architecture enhances performance and reduces storage requirements. When you enable Embedded Reports, you can navigate to a set of pre-built dashboards that chart multiple environment details. Dashboards and charts are powered by the Grafana observability platform. With Grafana, it's easy to navigate the existing dashboards, and to make your own charts and dashboards with no coding required.

For information about enabling Embedded Reports, see the *Installation Guide*.

- **Improved Metrics Discovery for vCenter Server Environments**

For vCenter Server environments, this release introduces an option to collect advanced Guest memory metrics. We introduce this new feature as a public preview.

When you specify a target for vCenter Server version 6.0 or later, you can turn on the **Enable Guest Metrics** option. Discovery will then access advanced metric data that the target monitors for its workloads. This can increase the accuracy of the VMEM data that Turbonomic uses for analysis of VMs.

Note that VMware Tools must be installed and running on the affected VMs. The **Hypervisor VMEM for Resize** option must also be turned on for the VMs discovered by the affected target. You make this setting in the Actions section of the automation policy for the given scope of VMs. This is turned on by default. For more information, see "Hypervisor VMEM for Resize" in the *User Guide*.

- **Reservation Improvements**

This release improves the performance when you create reservations. Turbonomic calculates placement more quickly, and it catches errors such as duplicate names for entities.

For API developers, the call to create a reservation is the same, except you do not pass the `apiCallBlock` parameter at creation time to make the reservation call synchronous. Instead, you create the reservation, and then execute a call to `GET` the reservation by ID, and pass the parameter `callBlocking=true`.

Configuration Requirements

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

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 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 fails to validate, this is 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:** This release includes policy settings to configure Response Time SLO and Transaction SLO. You can specify an upper limit for transactions or for response time. You can also set whether to turn SLO on or off.
- **Improvement:** Horizon VDI targets now report the number of connected sessions in use, as well as the total number of connected sessions (in use and not in use).

Fixed Issues

- **Fixed:** For large Kubernetes environments, it is possible that discovery will time out before it completes, resulting in incomplete discovery of the environment.
- **Customer Issue 111956**
Fixed: Upon upgrade to 7.22.4, the COST component continually crashes and fails to start up. This is because of a missing configuration property in the cr.yaml file.
- **Customer Issue 111916**
Fixed: After an update to version 7.22.4, RI analysis can be impacted. RI Utilization and Coverage figures can show a large drop, and analysis can stop generating Buy RI actions.
- **Customer Issue 111581**
Fixed: For Horizon VDI environments, Turbonomic does not include empty Active Directory groups as related groups in desktop pools. As a result, a move of a user to an affected pool can fail, or the user can move as an individual user and not as a member of a group.
- **Customer Issue 111668**
Fixed: Because you need the API to configure an Action Scripts target, this must be documented. The Target Configuration Guide should describe how to configure the target, and it should give a reference to the API guide for more information.
- **Customer Issue 111835**
Fixed: FOR PUBLIC CLOUD ENVIRONMENTS, You can create a policy that excludes specific VM templates for a given scope. In that case, Turbonomic should not create actions to resize the affected VMs to the excluded templates. However, analysis does not use the exclusion setting.
- **Customer Issue 111836**
Fixed: For Container environments, under some circumstances a plan can generate actions that impact entities outside of the plan's scope.
- **Customer Issue 111821**
Fixed: For a dynamic group that you create by combining filters with the OR operator, you cannot save a policy that uses that group as its scope.

- **Customer Issue 111802**
Fixed: Host IO in the user interface displays unusually low values. The values are calculated in GB/sec, but the user interface shows them as Gb/s.

- **Customer Issue 111804**
Fixed: For SaaS installations of Turbonomic, a user with the Site Administrator role cannot successfully change the password for a user account.

- **Customer Issue 111782**
Fixed: For UCS environments, under some circumstances discovery can fail when it encounters excessively large values for total throughput on a port.

- **Customer Issue 111767,111920**
Fixed: For scoped users, when the user displays a list of actions, the list can show the actions count, but it does not load the action entries into the list.

- **Customer Issue 111726**
Fixed: The Monthly Savings or Investments Totals chart has not been properly implemented, but the user interface enables you to add it to a dashboard. It gives erroneous values. This chart should be disabled.

- **Customer Issue 111456**
Fixed: For environments with large numbers of targets, under some circumstances the discovery process can reach a bottleneck where discovery seems to halt.

- **Fixed:** Performance is unacceptable when loading some charts or when loading the supply chain.

- **Customer Issue 111644**
Fixed: Under some circumstances action generation for Resize actions can appear to freeze. This can occur if you are using percentiles and a maximum observation period to gather associated utilization data. If you experience this issue, you should see a `Regular maintenance failed` error in the log for the topology-processor component.

- **Customer Issue 111411**
Fixed: When using filters to create a dynamic group, if you filter by cluster name, then the group can include VMs that are outside of the specified cluster.

- **Customer Issue 111319**

Fixed: For New Relic environments, under some circumstances discovery exceeds paging limits on the data it requests from the New Relic API. As a result, charts for scopes of managed databases or database servers can fail to display data.

- **Customer Issue 111686**

Fixed: When managing the list of policies, you can use a search filter to limit the list. Then, if you click use Select All, Turbonomic should only select the members that are in the filtered list. However, it selects the full, unfiltered set of policies in your installation.

- **Customer Issue 111644**

Fixed: For SaaS deployments of Turbonomic, performance to bring the installation up to date after a version upgrade is unacceptable.

- **Customer Issue 111267**

- **Customer Issue 111607**

Fixed: Under rare circumstance Turbonomic can recommend resize up actions for an environment that has underutilized VMs.

-

- **Customer Issue 111522**

Fixed: Under some circumstances, Turbonomic can fail to generate vCPU resize actions.

- **Customer Issue 111105,111565**

Fixed: When downloading data from the Pending Actions list (as CSV or PDF), if the list is very long the download can fail.

- **Fixed:** Reservations must now have a unique display name.

- **Fixed:** For Microsoft SQL Server environments, under some circumstances target validation can fail. This can occur when discovery encounters numbers that are formatted as strings.

- **Customer Issue 111409**

Fixed: Performance is not acceptable when you navigate to the Dashboards page, and the user interface loads the list of dashboards.

- **Customer Issue 110564**

Fixed: For Hyper-V environments, under some circumstances discovery can fail to discover a Hyper-V host. This can occur when discovery encounters DNS links to the host, and returns the link IP address.

- **Customer Issue 111448,111769**

Fixed: Under some circumstances, the Accepted Actions chart does not recognize all the actions that you have executed in your environment.

- **Fixed:** For cloud environments, under some circumstances the color of the supply chain token for a region does not reflect the severity of issues that impact that region.

- **Customer Issue 111430**

Fixed: Under some circumstances analysis attempts to execute a storage move to a disk that is already full. The underlying hypervisor rejects the action.

- **Customer Issue 110745,111375,111680**

Fixed: When running a plan, resize down actions for VMs do not reflect in the providers as reduced consumption of provider resources. For example, resize down of VMem does not show as reduced utilization of MEM on the underlying host.

Known Issues

- 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.

- **Customer Issue 105693**

The Headroom chart for All On-prem Hosts does not 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.

- Charts set to a time range that is greater than 2 hours do not display Cluster Statistics correctly.
To display data for time ranges of 24 hours or greater, Turbonomic must roll up the real-time data that it collects. However, it does not roll up data for cluster statistics. As a result, charts for clusters might not show valid data if they are set to these larger time ranges.
- Under some conditions, Cluster Headroom charts show Zero values for all their data points.
- For vCenter Server environments, VM moves across datacenters (from one vCenter to another) fail with the error, `Failed to operate with feature ActionExecution of probe class com.vmturbo.mediation.vmware.sdk.VimSdkConversionProbe.`
- For ServiceNow environments, Turbonomic fails to save any automation policy that sets the Action Type to **Request Approval from ServiceNow**.
- For ServiceNow environments, Turbonomic does not send any action to ServiceNow for auditing when the Action Type is set to **Generate action and send record to Service Now**.
- For AppDynamics environments, Turbonomic cannot discover databases if the target authentication uses `oAuth` for credentials.
- When calculating reservations in the PLACE page, if you select a Merge Policy as a placement constraint, then the placement will fail.
Note that if you have created a merge policy for a given cluster, and that cluster is in the scope of your placement request, then placement automatically uses that merge policy. You should not choose the merge policy as a placement constraint.
- The Place Page does not update after you calculate placement and create a reservation. To see the results of a reservation or placement request, you must manually update the page in the browser.
- 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.

- 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.

- For updates from versions earlier than 7.22.4, the update does not fully migrate policies for Application and Application Server entities.

Starting with version 7.22.4, the supply chain for applications has changed. Application and Application Server are now represented by the Application Component entity type. If you had created policies that affect these older entity types, then many of the settings will revert to their defaults. Before updating to the new version, you should check for affected policies and record the settings.

- When you create reservations (in the Placement page), if you provide a Network constraint the reservation does not recognize that constraint. The user interface displays a notification that the reservation cannot be created. However, the platform does create the reservation, and it does not include the network constraint.
- For vSAN environments, under certain circumstances a plan to add workloads can fail to place workloads, or it can fail to generate actions to increase storage capacity by provisioning new hosts.
 - If you scope the plan to a user-created group that only provides vSAN storage, or to a discovered storage cluster group, then the plan can fail to place VMs with multiple volumes. This can occur for VMs that use conventional storage (not vSAN) along with vSAN storage.
 - If you scope the plan to a vSAN host group and add workloads, the plan can fail to increase storage capacity by provisioning new hosts. For example, assume you scope the plan to a vSAN host group and add 20 VMs to the environment. In that case, you need hosts to provide compute capacity for the VMs, and you also need hosts to provide storage capacity. The plan can represent the compute provisioning correctly, but it can incorrectly fail to add more storage capacity to the vSAN.

- If the vSAN RAID type is `Raid6/FTT=2`, if you scope the plan to any vSAN groups then the plan will fail to place any of the VMs.
- Under certain circumstances an Optimized Cloud plan can fail to use newly purchased RIs, even if they match the plan environment and workloads. As a result, the costs that the plan shows can be incorrect.
- For Azure environments, when you inspect resource groups, Turbonomic does not currently show the billed costs for those resource groups.
- **Customer Issue 111396**

For cloud environments, under rare circumstances Turbonomic 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.
- The user interface includes a feature to configure email and trap notifications, and the User Guide includes a description of this feature. The user interface accepts and saves your configuration, but Turbonomic does not generate any notifications.
- After you update Turbonomic from the 7.21 version family up to the 7.22 version family, when you review saved plans the plan results do not include Storage Amount data. To regenerate the Storage Amount data, run the plans again.
- For public cloud environments that include AWS and Azure, when you run the Optimize Cloud plan with a scope that includes All Providers, the RI Coverage and RI Utilization charts do not display data for AWS. To view AWS data, scope the plan to only AWS providers.
- It is possible to set the Observation Period for Percentile utilization analysis to a value that is greater than the length of data retention for historical data. For example, if you set the observation period to 90 days, that is longer than the default 60 days of data retention.

To use a 90 day observation period for percentile analysis, be sure to increase your data retention to 90 days as well.

- **Customer Issue 110650**

Known Issue: When retrieving statistics for a cluster, using the `/groups` endpoint will return the aggregated statistics for all applicable cluster entities, and using the `/stats` endpoint returns the statistics for each individual entity.

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

```
STRICT_TRANS_TABLES,NO_ENGINE_SUBSTITUTION
```

In particular, the SQL modes should *not* include `ONLY_FULL_GROUP_BY`

- For Azure environments, Database resize actions do not properly consider storage capacity. As a result, Turbonomic can recommend resize down actions that are too aggressive, or it can fail to recommend appropriate resize actions. You should use `MANUAL` or `RECOMMEND` action modes, and verify that recommended actions are appropriate. Turbonomic is aware of this problem and is working to address it as soon as possible.

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

- For Azure environments with VMs in Scale Sets, for any VMs that are powered off the associated storage shows a utilization of zero GB. This is an accurate presentation of the data that the Azure environment returns for such a powered-off VM. However, it is likely that some of the storage capacity is currently utilized.

- **Customer Issue 110123**

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.

- 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
```

- Under rare circumstances, when the Turbonomic platform restarts it can fail to mount the platform storage. This occurs when the `heketi` pod does not start up correctly. Turbonomic uses `heketi` and `glusterfs` pods for storage, and when `heketi` fails to start it cannot mount the storage.

To recognize this situation, use the following command to monitor the `heketi` and `glusterfs` pods:

```
kubectl get pods -A | egrep "glusterfs|heketi"
```

You should see messages similar to the following:

```
Warning FailedMount 79s kubelet, node1 MountVolume.Setup failed for volume
"db" : mount failed: mount failed: exit status 1
```

If this occurs, delete the glusterfs pod with a command similar to the following, where {Unique_ID} is the ID of the glusterfs pod:

```
kubectl delete pod -n default gluster-{Unique_ID}
```

- When you update from 7.21.0 to this version, the update process sets your **Data Retention** setting back to the default values. If you have made custom data retention settings, you should reset them after you update.
- For public cloud environments, under some circumstances VMs or other workloads can appear in the user interface in an UNKNOWN state. However, there is no logging to indicate why the entity is in that state. This can occur when discovery succeeds, but the targeted infrastructure or cloud account cannot provide a known state for the workload. If this condition persists, you can contact your cloud vendor for more information about the workload.
- Updates from the 7.17 version family to the 7.21 version family cancel and delete any reservations that you have set up in the Placement view. If you require these reservations, then you should configure the reservations again in the updated version of Turbonomic
- If you are performing an *Online Update* from a 7.17 version of Turbonomic, then your update can inherit old configurations for the maximum MySQL connections. In large environments, or environments with many users of the same Turbonomic instance, this can result in the error, `error code [0]; Too many connections;`.

Your Turbonomic instance should be configured for a maximum of 151 connections. You can find this configuration in the following files on the Turbonomic instance:

- `/etc/my.cnf.d/server.cnf`
- `/opt/turbonomic/kubernetes/etc/my.cnf`

To correct this issue, change the configuration to allow 151 connections. The new setting should read `max_connections = 151`. After you make these changes, then restart the database. Either open a shell session as `root` or use `sudo`, and run the command, `systemctl restart mariadb`.

For assistance, contact Technical Support.

- When you update Turbonomic from the 7.17 version family to 7.21.0, you can lose any Accepted Action charts that you have included in your dashboards and views. The update changes these Accepted Action charts to All Actions charts.

To resolve this problem, edit the All Actions charts to change them to Accepted Actions charts.

- **Update Deletes Saved Plans**

Because of changes to Turbonomic plans, when you update from the 7.17 version family to the 7.21. version family, the update process deletes your saved plans.
- 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 108841**

In NetApp environments, the storage controller shows 100% utilization when there are no more disks in a `SPARE` state that the storage controller can utilize in an aggregate. This does not indicate that the storage controller has no capacity.
- 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.
- 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.
- The Turbonomic audit log tracks all communications with the platform via HTTPS. The log entries should include the IP address of the requesting client, as well as the user account. However, the log entries do not include the IP address of the originating client.
- 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.

- If you run the Alleviate Pressure plan in Turbonomic 7, and then compare it to the same plan and scope in a 6.1 release, then the display of instances in the supply chain are not identical for both versions.
- 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.
- For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.

- **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.
 - 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 you run a plan, the user interface enables you to make changes to the plan configuration and then run the plan again. If you do this, the plan results will be inconsistent. If you want to run a plan with a different configuration, you should start a new plan.
- You should never use duplicate names for groups of the same entity type. However, the user interface does not validate group names to keep you from creating a duplicate name.
- 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.
- For AWS environments, under very rare circumstances you can have RIs on payment plans that do not resolve to 1-year or 3-year terms. In this case, AWS does not return pricing data for those RIs. Turbonomic does not include such RIs in its calculations of RI utilization or RI cost.
- 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.
- **Customer Issue 109389**
In vCenter Server environments that have Instant Clone VMs, under some circumstances Turbonomic cannot move these VMs to other hosts in the cluster, even though you can manually migrate them via the vCenter Server user interface.