



---

# Turbonomic 7.21.5 Release Notes

April 27, 2020

This document describes issues that are addressed in Turbonomic 7.21.5 – Release Date: April 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](mailto:support@turbonomic.com), or open a ticket at:

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

## What's New for Version 7.21.5

This release of Turbonomic includes the following new features:

- User Interface Changes

For public cloud environments, the user interface now refers to *cloud tiers* instead of *templates*. In public cloud environments, you choose workload instance types to deploy among different tiers. For example, for VM instance types you choose among compute tiers, and for storage you choose among storage tiers. As you configure policies and review your cloud environment, the user interface refers to these as *cloud tiers*.

Turbonomic still uses templates for planning. For example, when you configure a plan you can choose templates for VMs to add to the plan. Turbonomic also creates templates to represent average utilization per cluster, and it discovers templates for VMs in your on-prem environment. You should remember that these templates are not the same as the cloud tiers that your public cloud services provide.

- Improvements for VMem Resize Actions

This release introduces the **Use Hypervisor VMEM for Resize** setting for VM policies.

For on-prem environments, Turbonomic discovers VMEM utilization and can recommend actions to resize the VMEM capacity on a VM. For environments that do not include any Guest OS Process targets, the data that analysis uses to make these recommendations comes from the underlying hypervisors. Unfortunately, that data is not

always sufficient to result in accurate resize recommendations. With this release, you can use the **Use Hypervisor VMEM for Resize** setting to determine how to generate VMEM recommendations:

- **On** (default)

When your environment includes Guest OS Process targets, Turbonomic uses the VMEM metrics those targets discover. If a scope of VMs does not fall under Guest OS Process targets, then analysis *will* generate VMEM resize actions for that scope. In this case, analysis uses the VMEM metrics it discovers from the underlying hypervisors.

- **Off**

When your environment includes Guest OS Process targets, Turbonomic uses the VMEM metrics those targets discover. If a scope of VMs does not fall under Guest OS Process targets, then analysis *will not* generate VMEM resize actions for that scope.

## 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 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:** We have changed the image names for the Turbonomic platform to increase their compatibility with general object naming restrictions.

# Fixed Issues

- **Customer Issue 110627**

**Fixed:** For API developers, the POST method to get stats does not always return the full set of data that you request. If you include a filter to return a count of entities (for example `numVMs`), along with other types of stats, the count of entities returns zero.

- **Customer Issue 110406**

**Fixed:** Under rare circumstances, the historical database can lock, which causes the user interface to freeze.

- **Customer Issue 110331**

**Fixed:** When working with schedules, if you select a schedule you can choose to defer the next occurrence. If you then edit the schedule (for example, change the times of occurrence), then the deferment should be lifted. However, the deferment remains, and the next *new* occurrence of the schedule will be deferred.

- **Customer Issue 110097**

**Fixed:** For Dynatrace environments, Turbonomic does not discover heap metrics for Tomcat servers.

- **Customer Issue 110395**

**Fixed:** For AWS environments, for a given AWS target Turbonomic discovers up to 50 Autoscaling Groups. If your environment includes more than 50, it does not discover them all. For the affected VMs, Turbonomic can recommend to resize the VM directly, instead of resizing the group configuration.

- **Customer Issue 109640**

**Fixed:** For API users, when you get lists of VM from the Search endpoint and the Targets endpoint, the Search endpoint can return VMs that have been deleted, while the Targets endpoint does not include the deleted VMs in its return.

- **Customer Issue 110397**

**Fixed:** The database maintenance action to rotate partitions requires privileges that are not necessarily appropriate for that user. As a result, if an enterprise changes the permissions for that user account, it can fail to execute the rotate partitions action. The action should not require those permissions.

- **Fixed:** When managing targets, you can select certain targets and delete them. If you use a filter to only display matching targets in the list, you can then choose **Select All** to select the filtered list of targets. However, that action actually selects all the targets you have defined. If you then delete the selection, you will delete all your targets.

- **Customer Issue 110386**

**Fixed:** For the Necessary Investments and Potential Savings charts, when you view them as text or bar charts, cost values do not include the currency.

- **Customer Issue 110371**

**Fixed:** When you create a group, you can set up filters to build lists of entities. The filters only execute the AND operator, they do not execute the OR or the XOR operator.

- **Customer Issue 110207**

**Fixed:** For API development, when you edit a workload placement policy via the PUT method, the API returns an empty `PolicyApiDto` object. The API should return an object that reflects the modified policy.

- **Customer Issue 110339**

**Fixed:** Under some circumstances, when you do not have connection to the internet, Turbonomic will not display schedules or save them to policies. This can happen if the schedule calls out to the internet to get updated timezone data.

- **Fixed:** For some environments, the topology processor can run out of memory. Turbonomic now recommends a minimum of 128 GB memory for the machine that hosts Turbonomic.
- **Fixed:** When you run an Optimize Cloud plan and view the RI Inventory chart, under some circumstances the *Show All* view can appear with no content. This can occur when you set the plan scope to a group.
- **Customer Issue 110334,110337,110347**

**Fixed:** Under some circumstances with very large environments, the topology-processor component can fail to start up because it is blocked by the history pod. When it is in this state, it continually tries to restart.
- **Customer Issue 110095**

**Fixed:** In VSAN environments, multiple hosts can provide storage resources for a single VM. In the supply chain, this can appear as though a single VM is hosted by multiple physical machines. This is confusing. The supply chain should only show the primary host that provides compute resources.
- **Fixed:** For public cloud environments with billing families, if you set the scope to a billing family, the user interface displays a count of workloads for the billing family master account. The master account should not show that it contains workloads. Instead, the billing family member accounts should be the ones to show counts of workloads.
- **Customer Issue 110015**

**Fixed:** When working with a placement policy, you can use the **Enabled** option to either enable or disable that policy. Discovered placement policies (for example, for DRS rules), should be read-only, and you should not be able to enable/disable the policy or edit and save your changes. However, the user interface does not disable these controls. This results in a confusing user experience.
- **Customer Issue 110011**

**Fixed:** Under some circumstances, the Health chart display applies the wrong set of issues to the given color. For example, assume 10 Without Risks, 20 with Minor Risks, and 5 with Critical Risks. The chart could show 20 green, 10 yellow, and 5 red.
- **Customer Issue 108355**

**Fixed:** When you create a group of entities, you should be able to navigate to the Search view, and search for that group. However, search does not return your created group.
- **Customer Issue 102278,104021,107368**

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

## Known Issues

- Under certain circumstances, an update from a version earlier than 7.21.5 can cause schedules to be removed from scheduled policies.

This can happen if you created a schedule with the *z* timezone. This is not a valid timezone, and version 7.21.5 no longer supports it. After you update, you should check your scheduled policies to ensure the schedules are still attached.
- Under some circumstances, the Headroom charts can fail to show complete data. When this occurs, the charts can fail to show data for a 2-hour time frame or for a 1-year time frame. In addition, the charts can fail to show data for all the resources they monitor.

- For **Merge** placement policies, if you merge datacenters (as opposed to merging clusters) then the policy can fail to recognize the merge.

For example, assume you merge Datacenter A and Datacenter B in a policy, and circumstances arise where a VM is forced move from Datacenter A to Datacenter B. In that case, Turbonomic can fail to create the move action. Instead, it can create a Reconfigure VM action that would make the VM comply with the placement on Datacenter B.

For other circumstances where a VM move across datacenters is optional, Turbonomic will not generate the action.

- For public cloud environments, when you configure the Optimize Cloud Plan, the RI Inventory does not appear in the configuration display.
- For environments with Microsoft EA targets, under some circumstances EA discovery times out and fails to complete. The log includes messages similar to:

```
log.txt_20200415_174502:mediation-azureea-c8448796b-wz2r Caused by: java.net.SocketTimeoutException: Read timed out
```

If EA discovery fails to complete and you see this error, contact your support representative for assistance.

- 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.
- For Azure environments, under some circumstances you might need to manually start the analysis process to calculate Azure RI Buy actions.

Upon startup, because of timing issues for discovery of Azure billing entities (for example, Billing Families or costs), Azure RI analysis might not execute. Also note that if you restart the Turbonomic Action Orchestrator component for some reason, then Turbonomic can lose the currently pending Buy RI actions. Manually starting RI Buy analysis will correct these problems.

To manually force a restart of the RI Buy analysis, change the term of your Azure Profile settings. For example, navigate to **Settings: Billing and Costs: Reserved Instance Settings** and change the Azure Profile settings from **1 Year** to **3 Years**. This change triggers the RI Buy analysis.

- The All Actions chart does not include pending actions for databases or database servers.
- For public cloud environments, when you run the Optimize Cloud Plan, the results can show more RI Inventory than the plan's scope should allow. For example, if you scope the plan to a region that includes only a limited number of RI templates, the plan results will include all the RI templates in your entire environment.
- 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.

- 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.
- 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.
- 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.

- For AWS environments that use dedicated tenancy, under some circumstances Turbonomic can recommend changing a workload to a valid instance type for the tenant, even though the current type is already valid. This can happen when the instance type is not included in the Offer File for the tenancy.

For example, assume the `t3a` template family does not support dedicated tenancy. However, assume that the user created a `t3a` instance with dedicated tenancy in the EC2 console. In that case, Turbonomic will see this as a misconfiguration and recommend changing to a different instance type.

- 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.
- 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 spot instances of VMs, the Turbonomic does not discover price data for the VM instances. The user interface does not show pricing, and the pricing is not available for analysis.
- 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 specify a schedule, you must be aware of issues for Daylight Savings Time changes. For example, assume you create two schedules, one that starts at 7:00 am and lasts 12 hours, and another schedule that starts at 7:00 pm and lasts 12 hours. If you apply those schedules to policies on the same scope, then the policies can experience a scheduling conflict as Daylight Savings time changes.
- 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.