



Turbonomic 7.21.1 Release Notes

March 3, 2020

This document describes issues that are addressed in Turbonomic 7.21.1 – Release Date: March 3, 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>

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.

Fixed Issues

- **Fixed:** For the REST API, when you request paginated data the API returns a header which indicates to the caller what Cursor to request in order to get the next page. For `markets/entities`, the header does not properly indicate the cursor to request.
- **Customer Issue 109532**
Fixed: When you increase the retention settings for hourly data, you should be able to see hourly data in charts for a longer historical period. However, these settings do not have an effect on chart display.
- **Customer Issue 109168**
Fixed: When you configure a Hardware Replace plan, the Host Utilization setting has no effect.

Known Issues

- When you update from 7.21.0 to version 7.21.1, 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 execute actions in Turbonomic that result in a savings, different charts in the user interface should show these savings. However, some charts do not show those results. For example, the Top Accounts chart does not show savings in the Saved By Actions and Actions Taken columns.
- 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.
- For public cloud environments that use New Relic, at a maximum Turbonomic can discover 2000 VMs that host applications managed by New Relic.
- 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.
- When you review headroom statistics, the Memory Headroom, Storage Headroom, and CPU Headroom charts do not show the VM template that was used to generate the chart data.
- 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
- For environments that include multiple billing families, it is possible that Turbonomic can assign resources from the wrong RI to a workload. If the workload is in an account that belongs to Billing Family A, Turbonomic can incorrectly assign an RI from Billing Family B to the workload.

- Possible Stranded Data for Deleted Plans

When updating from the 7.17 version family to the 7.21 version family, Turbonomic deletes your saved plans. For more information, see [Update Deletes Saved Plans \(on page 4\)](#).

However, under some circumstance it removes the plans from the user interface, but does not fully delete the plan data. This can happen when some components fail to start up before the update deletes these plans.

When the update encounters this problem, it posts an error message to the log, similar to:

```
ERROR [main] [V_01_00_00__PURGE_ALL_LEGACY_PLANS] : org.apache.http.conn.HttpHostC
onnectException: Connect to arangodb:8529 [arangodb/10.233.6.158] failed: Connectio
n refused (Connection refused)
```

In future releases of Turbonomic, the update process will identify this stranded plan data and delete it. If you suspect the stranded data is consuming too much space, please contact Technical Support for help deleting this data. For example, if you have a very large environment, and you saved a large number of plans, you might consider deleting this data.

- 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 `readmax_connections = 151`. After your 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.

- 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 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 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 environments that use Raw Device Mapping (RDM) storage, it is possible for the user interface to show storage counts that seem inconsistent. This is because a VM can have all the potential RDM blocks as providers, while a Host has only the specifically used RDM blocks as providers.

For example, when you scope the view to the VMs on a host, the supply chain can show a higher count for storage than when you scope the view to the given host. Or when you run a plan to merge clusters, the plan results can include storage that is outside of the merged clusters.

Note that this only affects display in the user interface. Turbonomic analysis in real-time and in plans uses the correct storage relationships.

- 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.

Under some circumstances, the update does not fully delete the plan data. For more information, see [Possible Stranded Data for Deleted Plans \(on page 3\)](#).

- 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.

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

RI Savings for Individual Accounts

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

Disabling Resize for VM Resources

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

- For AWS environments, if a target account includes a policy with region-level control, then discovery of that target will fail with an `authorization` exception.

- **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.

- 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.