



---

# Turbonomic 7.17.8 Release Notes

December 20, 2019

This document describes issues that are addressed in Turbonomic 7.17.8 – Release Date: December 20, 2019. 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>

## 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:** When installing Turbonomic in your Kubernetes environment, you now have the option to provide a secret key to access the product images. If you prefer, you can still provide username and password.

# Fixed Issues

- **Fixed:** For installations that use an external database (not the database component that ships with Turbonomic), if the database is not synchronized to the UTC time zone, then some procedures to manage historical data can result in loss of data.

**NOTE:**

The *Installation Guide* now lists time synchronization as a required step. It states that you must synchronize the Turbonomic platform to the UTC time zone.

- **Fixed** As Turbonomic performs rollups of historical data (aggregating data points for hourly, daily, weekly, etc. time scales), under some circumstances it can omit data from the rollups.
- **Fixed:** Under some circumstances, for a specific type of queries the user interface sends to the database, it is possible for the query to fail. As a result, user interface components can fail to show data.
- **Fixed:** When the Turbonomic database is rolling up historical data (into hour, day, week, etc. aggregations), there are circumstances where time synchronization issues can result in a loss of data. As a result, charts in the user interface can display with no data.

For example, this can occur if you have issues with NTP configuration. The database rollup procedures should be more robust when such issues arise.

- **Fixed:** Under some circumstances, after removing targets, Turbonomic does not remove clusters or other groups that were created via discovery through those targets.
- **Customer Issue 108985**

**Fixed:** When creating a plan to add workload, you can specify maximum utilization for resources in the plan scope. However, if the plan is scoped to a cluster, that setting does not take effect, and the plan can exceed that max utilization. However, the setting does take effect if you scope the plan to a custom group that contains the same hosts.

- **Customer Issue 108907**

**Fixed:** Under some circumstances, as you change the scope to different clusters, scoping to a given cluster results in the Optimized Improvements chart displaying no data.

- **Customer Issue 108928**

**Fixed:** In large environments, some navigation gestures (for example, navigating via the supply chain) can overload the data payload passed between the user interface and the platform. This can result in the following error:

```
RESOURCE_EXHAUSTED: gRPC message exceeds maximum size.
```

- **Fixed:** The performance for some database queries from the user interface is unacceptable. As a result, the user interface can appear sluggish or frozen.
- **Fixed:** For logging into the Turbonomic user interface from a Macintosh system, if you upgrade your OS to Catalina 10.15, then you cannot log in using a Google Chrome browser.

# Known Issues

- **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.

- **WARNING:** Do Not Delete the `administrator` User Account.

The user interface allows an administrator user to delete all the user accounts for a Turbonomic installation. This includes the `administrator` user account. *You should NEVER delete all user accounts, and you should never delete the administrator account.*

- 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.