



Turbonomic 7.22.3 Release Notes

June 26, 2020

This document describes issues that are addressed in Turbonomic 7.22.3 – Release Date: June 26, 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 7.22.3

This release of Turbonomic includes the following new features:

- VCPU and VMem Utilization Charts for VMs

Turbonomic uses percentile calculations to measure VCPU and VMem utilization more accurately, and drive scaling actions that improve overall utilization and reduce cost for cloud VMs. When you examine the details for a pending scaling action on a VM, you will now see charts that highlight VCPU and VMem *utilization percentiles* for a given observation period, and the projected percentiles after you execute the action. The charts also plot *daily average utilization* for your reference. If you have previously executed scaling actions on the VM, you can see the resulting improvements in daily average utilization. Put together, these charts allow you to easily recognize utilization trends that drive Turbonomic's scaling recommendations.

For on-prem VMs, you will see either a VCPU or VMem chart, depending on the commodity that needs to scale. For cloud VMs, both charts display.

These charts also appear when you scope to a given VM (on-prem or cloud) and view the Details page.

- Improved Scaling Constraint for Cloud Workloads

This release introduces the **Cloud Instance Types** scaling constraint in policies for VMs, databases, and database servers. This is an improvement on the **Excluded Cloud Tiers** setting in previous versions. With this change, you now select instance types for the scoped workloads and see their resource allocations right away.

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:** For deployments that use Single Sign On (SSO), with this release Turbonomic can authenticate a user when SAML returns more than one user group for the given credentials. In the case of multiple groups, Turbonomic authenticates the user as a member of the group with the least privileges.

Fixed Issues

- **Customer Issue 110759**
Fixed: For Azure environments, under some circumstances with a large environment, timing issues cause analysis to use default policy settings even though custom policies are set for the given scope of entities.
- **Customer Issue 111353**
Fixed: For AWS environments, if your environment does not have access to specific AWS regions, Turbonomic should discover that and only try to discover the regions you can access. However, under some circumstances discovery can fail as Turbonomic tries to discover regions that your target cannot access.
- **Fixed:** For installations that use a remote database, under some circumstances updates of the modeled environment can show poor performance. This impacts the performance of analysis.
- **Customer Issue 111371**

Fixed: For some environments, Turbonomic includes "proxy" entities in its count of managed workloads. The count goes against the license capacity.

For example, in a DataDog environment Turbonomic discovers entities through the DataDog target that map to entities discovered through other targets. These are "proxy" entities. When counting workloads to determine utilization of the license capacity, should not include these proxy entities.

- **Customer Issue 111356**

Fixed: When you scope Turbonomic to a public cloud account, you can navigate to a list of workloads for that account. The user interface includes controls to sort the list by Utilization, Name, and Severity. However, sorting by Utilization has no effect.

- **Customer Issue 111319**

Fixed: For New Relic environments, under some circumstances the charts for scopes of managed databases or database servers can fail to display data.

- **Fixed:** Under certain circumstances, the Storage Summary chart can show incorrect Potential Savings. This can happen when the view scope is set to a Billing Family. To see the correct potential storage savings for a billing family, set the scope to the Volume nodes for the given Billing Family.

- **Customer Issue 111481**

Fixed: Under rare circumstances the running of full discovery and the running of incremental discovery can overlap. In this case both discoveries fail to complete.

- **Customer Issue 111124,111169**

Fixed: For application management environments (for example, DynaTrace), when you set the scope to Business Applications for Resource charts and Multiple Resource charts, the charts only show data for the Last 2 Hours timeframe.

- **Fixed:** For the Potential Savings chart, when you click **Show All** to display the Potential Savings table, and then scope it to Delete Volumes, the list displays unattached volumes that you can delete. However, for Azure environments this list can also include unattached volumes that are not eligible for deletion.

- **Customer Issue 110888**

Fixed For environments with empty clusters (no hosts in the cluster), the display of Headroom charts can be misleading. Turbonomic should not calculate headroom for empty clusters.

- **Customer Issue 110660**

Fixed: For Azure environments, if you have storage volumes that are not attached to any workloads, the charts for those volumes do not show IOPS or storage capacity.

- **Customer Issue 111251**

Fixed: For very large environments, under some circumstances when loading the topology Turbonomic can experience a `java.lang.OutOfMemoryError: exception`.

- **Customer Issue 111090**

Fixed: For the Top N Virtual Machines chart, when you download the chart data as CSV, the data does not include a column for "Storage" data.

- **Customer Issue 111036**

Fixed: For the RI Inventory chart, when you download a CSV file from the chart, the data shows Account ID, but it should show Account Name.

- **Customer Issue 111156**

Fixed: For very large environments, timing issues can result in intermittent gaps in the ingestion of incremental discovery data.

- **Customer Issue 110928**

Fixed: For hybrid environments, if you create a resource group that includes cloud and on-prem entities, then charts that are scoped to that group will not show cost data. This is true even for the cloud entities in that scope. The charts should show cost data for the cloud entities.

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

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.

- For Turbonomic 7.22.3, an issue with SAML redirects prevents Single Sign On (SSO) from authenticating user accounts.

If you are installing Turbonomic for the first time, do not enable SSO until a fix becomes available.

If you are updating to this release, and your instance has already enabled SSO, then you must disable SSO until a fix is available. For instructions to disable SSO, see "Disabling Single Sign-On" in the *Turbonomic Installation Guide*.

- If your instance of Turbonomic is inoperative for a period of time, that can affect the cost calculations for Migrate to Cloud Plans.

To calculate costs for a VM that it will migrate to the cloud, Turbonomic considers the VM's history. If the VM has been stable for 16 of the last 21 days, then Turbonomic will plan for that VM to use an RI. In this way, the plan calculates the best cost for the migration. However, if Turbonomic is inoperative for any time, that can impact the historical data such that the plan will not recognize a VM as stable even though it is.

- **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.
- Plans to add host entities by template fail if you use HCI Host templates.

When you create a plan, you can choose to add new hosts to an environment, or you can choose to replace existing hosts with a different host type. One way to specify the new hosts is to use Host templates. If you use an HCI Host template, then the plan will fail. If you use normal Host templates, the plan can succeed.

- An update from a version earlier than 7.22.3 can affect the display of historical data in charts and action details. The affected historical data does not affect analysis or the actions Turbonomic generates.

For charts that are scoped to cloud VMs, the display of Storage Access, Storage Amount, and IO Throughput commodities can show a break in historical data.

The action details for Delete Volume actions (for cloud volumes) can show "No Data" in the fields for Last Attached VM and Time Unattached. This will persist until Turbonomic discovers new data for these fields.

- **Customer Issue 111370**

When you configure an AWS target to use an IAM role for credentials, Turbonomic fails to validate the target.

AWS target configuration includes an option to authenticate via Access Key and Secret Access Key, or via IAM Role. When you choose IAM Role, validation fails because it incorrectly expects a value for the Secret Access Key. To work around this issue and configure an AWS target via IAM Role:

1. Enter an arbitrary string for the Secret Access Key.
Turn off **IAM Role** and enter an arbitrary string in the **Secret Access Key** field.
 2. Configure your IAM Role.
Turn on **IAM Role** and enter your role in the **IAM Role ARN** field.
 3. Complete the target configuration and click **ADD**.
- 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.

- 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.

- **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.

- Known Issue:** When using the POST /stats request, including the `numStorages` statistic in the filter causes some other statistics to not appear, resulting in a partial response.

Do not use the `numStorages` statistic filter. This statistic *will* be returned when other filters or no filters are used.

- 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.
- 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 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.
- 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 anti-affinity 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.