



Turbonomic 7.22.2 Release Notes

June 10, 2020

This document describes issues that are addressed in Turbonomic 7.22.2 – Release Date: June 10, 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 7.22.2

This release of Turbonomic includes the following new features:

- Feedback and Diagnostics

Turbonomic constantly strives to improve your experience with our Application Resource Management platform. In an effort to better understand how well the user interface responds to your requests and interaction, we have introduced **Feedback and Diagnostics**. This is an Opt-In feature that collects anonymized usage data and environment statistics. With this data, we can see how responsive the platform is.

For example, when you open the Plan page in the user interface, Turbonomic records the page load event, and collects non-confidential data such how long it took to load the page. This information helps us identify and prioritize areas that we can improve.

Only users with the **Administrator** role can enable or disable this feature. Turbonomic never uses this feature to collect sensitive data such as IP addresses or geographical locations.

- Consistent Resize

For certain groups of entities, Turbonomic discovers that the entities should use consistent resizing, and so it creates consistent scaling groups for them. In earlier versions Turbonomic set consistent scaling internally, and there was no visible representation of that setting for the user.

Turbonomic discovers such groups and creates read-only policies for them to implement their consistent resizing. You can see the policies assigned to the groups. To disable consistent resizing, you create another policy for that scope and turn off **Consistent Resizing**.

- Improvements for Potential Savings or Investments Charts

These charts show you the total savings or investments that Turbonomic recommends for the given scope. But how do you inspect the individual actions that will give you these savings? With this release we introduce the **Show All** button to open a tabular view with rich details about the given investments or savings. This tabular view:

- Lists all the associated actions, sorted with the biggest investment or savings impact first
- Categorizes the list by action type and entity type for easy filtering and navigation
- Shows the before and after savings you gain from RI coverage, for individual workload instances
- Includes details so you can fully inspect each action

- Improved Container Management

For Kubernetes environments, this release introduces automation policy settings that you can use to fine tune actions for Container and Container Spec entities:

- Container:

This release introduces Increment Constants for container VCPU and VMEM, as well as a Low, Medium, or High Rate of Resize. This release also introduces Consistent Resizing for container groups. For example, Turbonomic turns on Consistent Resizing for container groups as a way to adhere to Container Spec settings.

- Container Spec:

For Container Specs you can set the percentile *aggressiveness* when analyzing VCPU and VMEM utilization. Along with this, you can set the min and max observation periods to calculate the percentile utilization.

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

- **Customer Issue 110523**
Improvement: With this release we have made the *Loading* spinner icon easier to see in the user interface.
- **Customer Issue 110852**
Improvement: The Potential Savings and Necessary Investments charts now include a **Show All** button to provide tabular data for the entities that the given chart tracks.

Fixed Issues

- **Fixed:** For the vCenter Server probe, to address security vulnerabilities the probe must update the version of `jackson-databind` library that it uses.
- **Customer Issue 111202**
Fixed: For the configuration of Active Directory authentication, the Installation Guide does not specify that the trust store must be of a JKS type. The command to configure the store should be: `keytool -import -alias ldapcert1 -file /tmp/ldapserver.pem -keystore cacerts -deststoretype jks -storepass changeit -noprompt`.
- **Customer Issue 111095**
Fixed: For Azure environments, in cases where Turbonomic does not have pricing information for a given VM instance type, under some circumstances it can recommend resizing to that type as an efficiency action, even though the instance type is more expensive.
- **Customer Issue 111069**
Fixed: When configuring a Datadog target, Turbonomic does not validate the Proxy Port before trying to validate the target. If you provide an invalid Proxy Port entry then Turbonomic will use that to attempt a connection. This can result in the target failing to validate, sometimes with a viable error statement, and sometimes with no error statement.
- **Customer Issue 111080**
Fixed: When running plans with **Ignore Constraints** turned on, there is a rare circumstance where the plan can fail with a stack overflow error. If you set the scope to include a group of groups that includes itself as a member, then **Ignore Constraints** can cause a plan for that scope to fail.
- **Customer Issue 111325**
Fixed: For Azure environments when running an Optimize Cloud Plan, under some circumstances the analysis fails to get pricing information for entities, and the plan fails to complete.
- **Customer Issue 111032**

Fixed: Under some circumstances Turbonomic discovers networks that the given workloads cannot access. In that case, it can recommend reconfiguring VMs to use networks that they cannot access.

- **Customer Issue 110993,111013**

Fixed: Under certain circumstances, Turbonomic fails to discover Azure pricing data, and logs errors.

- **Fixed:** Turbonomic generates average utilization templates to use when it analyzes cluster headroom. For environments that have different datacenters, if those datacenters have clusters with the same names, then the generated templates would also have the same name. In that case, Turbonomic only generates a single template of that name.

- **Customer Issue 110970**

Fixed: For actions such as Resize Up for database server capacity, the action details and action log show decimal values for MB or GB, when they should round to the nearest integer.

- **Customer Issue 110973**

Fixed: For Hyper-V environments that include Static Memory VMs, under some circumstances the discovered memory capacity can be incorrect.

- **Customer Issue 110857**

Fixed: For public cloud environments with VMs that do not enable memory monitoring, under some circumstances for a resize action the action details do not show the correct projected VMem capacity for the recommended instance type.

- **Customer Issue 110833**

Fixed: The Location chart that maps out the locations of cloud accounts and subscriptions does not include a legend.

- **Customer Issue 110796**

Fixed: For environments with business applications, the user interface lists the same business applications whether showing the Hybrid View, the On-Prem View or the Cloud View. The different views should filter the display of these entities accordingly.

- **Customer Issue 110132**

Fixed: For some VM Resize Down actions, Turbonomic powers down the VM. Under some circumstances, if the action fails then Turbonomic does not power the VM back up again.

- **Customer Issue 110590**

Fixed: The user interface includes a breadcrumb field to show the context of your session scope. If you set the scope to a Cloud Account, or to a Cloud Region, this field shows incorrect data.

Known Issues

- **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**.

- For VM actions, Turbonomic includes the **Enforce Non Disruptive Mode** setting. The documentation for this setting is incorrect.

The documentation states that the setting will enforce that any *disruptive* action will automatically demote action mode from MANUAL or AUTOMATED, down to RECOMMENDED. In fact, this setting demotes AUTOMATED down to MANUAL, and it does not demote MANUAL down to RECOMMENDED.

- 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.
- 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.
- 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 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.
- 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.

- 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.

- **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.

- 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 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.