



Turbonomic 7.22.0 Release Notes

May 13, 2020

This document describes issues that are addressed in Turbonomic 7.22.0 – Release Date: May 13, 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 7.22.0

This release of Turbonomic includes the following new features:

- Support for Azure Managed Ultra Disk Storage

For Azure environments, Turbonomic now discovers Managed Ultra Disk storage, and the respective costs for IOPS and Throughput. As part of discovery, Turbonomic recognizes the requirements that VM types must meet in order to use Managed Ultra storage. For actions to scale a VM, Turbonomic only recommends changing to a VM type that can support the indicated storage. This means you can automate VM scale actions, and the resized VMs will still use the appropriate Managed Ultra Disk storage.

- Shared-Nothing VM Migrations

This release introduces the **Shared-Nothing Migration** setting for actions in VM Automation Policies. To enable Shared-Nothing moves:

- Set the action modes to be the same for VMs and Storage in a given scope. The modes for both must be either AUTOMATED or MANUAL.
- Explicitly turn on **Shared-Nothing Migration** for that given scope of VMs.

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 109927**

Improvement: Starting with this release, you can copy a link to a specific dashboard or view in Turbonomic, and open it directly. Assume you have not logged in already - When you open that link you first log in, and then Turbonomic takes you to the given view. You can save bookmarks to views, or send links to other members of your team.

Also, when you log out of Turbonomic and then log in again, you log back into the same location that you last visited before you logged out. For example, if you were inspecting a specific scope of entities when you logged out, you can log back in directly to that same scope of entities.

- **Customer Issue 108434,108988**

Improvement: When searching for entities or creating groups, you can create a filter by Resource Group. With this release, you can now use REGEX expressions for the filter criteria.

Fixed Issues

- **Customer Issue 110799**

Fixed: For large Azure environments, under some circumstances the Azure EA probe can time out before it completes discovery.

- **Customer Issue 110907,110910,110920**

Fixed: For actions to move Virtual Volumes, the default action mode is Recommend. If you change the mode for a scope of volumes, under some circumstances Turbonomic ignores that new action mode. For example, if you change to Manual for a scope of volumes, Turbonomic still generates Recommend actions for that scope.

- **Fixed:** 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.
- **Fixed:** 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 to 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.

- **Fixed:** For public cloud environments, when you configure the Optimize Cloud Plan, the RI Inventory does not appear in the configuration display.
 - **Customer Issue 110589,110722**
- Fixed:** When using filters to create groups, it is possible to specify invalid REGEX for the filter. The user interface does not validate the REGEX string, and creates the group. However, when Turbonomic tries to load the group data, it fails with a SQL exception. Also, this bad group keeps other groups from displaying.

- **Customer Issue 110576**

Fixed: For Azure environments that include GPU-optimized VMs, under some circumstances Turbonomic can generate incorrect Reconfigure actions for the VMs. Assume a VM is based on a promotional template (for example, Standard_NC6_Promo). When that promotion expires, then you should move the workload to a different VM type. However, analysis generates a reconfigure action instead. Turbonomic does not support resizing GPU-optimized VMs, so it should not generate an action for this circumstance.

- **Customer Issue 110545**

Fixed: When viewing the Capacity and Usage chart, if you include Virtual Storage as one of the metrics to track, then the chart only shows Virtual Storage statistics.

- **Customer Issue 110556**

Fixed: In large public cloud environments, under some circumstances the details for a Delete Volume action do not show complete data.

- **Customer Issue 110405**

Fixed: For AppDynamics environment, under rare circumstances Turbonomic cannot validate the target, even though it can fully discover the target's environment. This can happen when the target does not provide connections for all of the endpoints that Turbonomic requests. In such circumstances, Turbonomic should post a warning.

- **Customer Issue 110404**

Fixed: For UCS environments, Turbonomic does not recognize breakout ports in IO Modules. As a result, the Fabric Interconnect might not appear as connected to the corresponding IO Module, and Turbonomic might not discover traffic over these ports.

- **Customer Issue 110436**

Fixed: For the SaaS version, Turbonomic must improve the performance when loading the user interface after you log in.

- **Customer Issue 110431**

Fixed: For VMware Horizon environments that use LDAP authentication for users, if the environment includes more than 1000 VDI users, then it is possible that Turbonomic can fail to discover the associated Business User entities.

- **Customer Issue 110297**

Fixed: Under rare circumstances, after you delete targets and add other targets of the same type, it is possible that this can generate duplicate data in your environment. This can cause components to fail, or the API can fail to return requested data.

- **Fixed:** For Hyper-V environments that include WMI targets, if target discovers dynamic memory via WMI, then it can report excessively high memory utilization values.

- **Customer Issue 109359,110319,110820**

Fixed: For UCS environments, under some circumstances the user interface includes chassis IDs when it displays cluster names.

- **Customer Issue 110288**

Fixed: Under rare circumstances, if the connection times out between a probe component and the Turbonomic topology component, then it can fail to establish the connection again. The related targets display in the user interface as invalid, and they do not reconnect.

- **Fixed:** 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.

Known Issues

- The Shared-Nothing Migration feature is not fully supported at this time.

This release adds a new policy setting to turn on Shared-Nothing Migrations for a given scope of VMs. The What's New section of the User Guide describes this feature. However, this feature is not fully supported yet. In most cases you should not use the feature. If you want to experiment with Shared-Nothing Migration, please contact your support representative.

- When you display the Turbonomic user interface in Internet Explorer 11 or earlier, some of the navigation icons do not display correctly. The interface displays their alternative text, and you can click them to execute the links.
- 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.

- 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 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.