



Turbonomic 7.22.7 Release Notes

August 21, 2020

This document describes issues that are addressed in Turbonomic 7.22.7 – Release Date: August 21, 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 7.22.7

Version 7.22.7

- **New Application Targets**

This release introduces support for the following targets:

- Apache Tomcat, versions 7.x, 8.x, and 8.5.x
- MySQL, all 5.6 and 5.7 releases

- **Consolidated Resize Actions for Containers**

Executing several container resize actions can be very disruptive since pods need to restart with each resize. To minimize disruptions, Turbonomic now consolidates resize actions into one merged action for replicas of the container scale group(s) related to a single workload controller. When a single action group has been executed, all resizes for all related container specifications will be changed at the same time, and pods will restart once.

To view and manually execute an action group, set the scope to the workload controller for the containers or container spec being resized. Merged actions are the default configuration for KubeTurbo version 7.22.7 and later.

- **User Interface Links**

When you right-click a link in the Turbonomic user interface, you now have the option of opening it in a new browser tab or window.

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 is 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:**

Starting with 7.22.7, it is easier to specify the memory capacity you want to assign to the platform upon installation. To set the VM memory to 128 GB or 64 GB, you no longer need to make changes to the platform's cr.yaml file. The installed platform automatically recognizes and uses the VM capacity that you set.

Turbonomic installations now take advantage of more robust memory management to keep the individual components running under more varied conditions.

- **Improvement:**

Customer Issue 112127

When you display an All Actions chart in Tabular Format, the table now includes additional columns, to give you richer details about each action in the list. When you export the table data, it also includes these additional fields.

Fixed Issues

- **Fixed Issue:**

Customer Issue 112337

Policies that set their scope to groups created with the `discoveredBy` filter do not take effect.

- **Fixed Issue:**

Customer Issue 112245

In Kubernetes environments that use multiple consistent scaling groups with different overall capacities, under some circumstances Turbonomic can recommend scaling actions that are appropriate for one scaling group, but apply them to a different scaling group. This can manifest as actions to scale a container down to a capacity that leaves the container highly overutilized.

- **Fixed Issue:**

For AppDynamics environments, under some circumstances can discover business services or transactions that depend on tiers that Turbonomic cannot discover, where the inability to discover is because of permissions. In that case, discovery can fail.

- **Fixed Issue:**

Customer Issue 111668

For Action Script configuration, the *User Guide* has an error in its description of the Manifest file. ActionScript examples describe configuration of the given action phase, `actionPhase: PREP`. This naming is incorrect for that action phase. It should be `actionPhase: PRE`.

- **Fixed Issue:**

Customer Issue 112203

Earlier versions included a problem in the user interface that prevented users from changing some policy values. This has been fixed in subsequent versions.

- **Fixed Issue:**

Customer Issue 112218

For Dynatrace environments, under some circumstances the request to the target for metric data exceeds the allowed length for a URI. In that case, discovery can fail to get the metrics it needs for the given entity.

- **Fixed Issue:**

Customer Issue 112186,112331

For Azure environments, the data files for instance types must be updated to include newly added AMD instance type families.

- **Fixed Issue:**

Customer Issue 112059

For Hyper-V environments, under some circumstances Turbonomic discovers datastores with the incorrect external name. Names should also include mountpoint when appropriate.

- **Fixed Issue:**

Customer Issue 112108

For Azure environments that set quotas on regions, if multiple targets share the same region, the Turbonomic can use the incorrect quota when calculating whether to generate an action for an entity in the affected region. As a result, it can generate actions that fail due to exceeded quota.

- **Fixed Issue:**

Customer Issue 112058

For Hyper-V environments, under some circumstances the naming for discovered clusters results in strings that are too long to store. As a result, Turbonomic might not display all the groups that are managed by a Hyper-V target.

- **Fixed Issue:**

Customer Issue 112103

For deployments that use an external MySQL database, under some circumstances can fail to recommend RI Buy actions.

- **Fixed Issue:**

After you update Turbonomic to a new version, under some circumstances when you refresh the browser the user interface stays in the "Initializing" state. If you experience this issue, then clear the browser cache and try to reload the application again.

- **Fixed Issue:**

Customer Issue 112048

Under some circumstances, Turbonomic can recommend unnecessary VM moves due to net throughput congestion.

- **Fixed Issue:**

Customer Issue 111998

For VSAN environments, under some circumstances Turbonomic can see storage that is attached to a VM as wasted storage.

- **Fixed Issue:**

Customer Issue 111829

For environments that use placement policies to *limit* placement (do not place together, etc.), an action that is generated to comply with the policy can be described as a Performance Assurance action instead of a Compliance action.

- **Fixed Issue:**

Customer Issue 111941

For large AWS environments, it's possible that the volume of data requests can result in discovery failing to collect all the metrics for the environment.

- **Fixed Issue:**

Customer Issue 111790

For vCenter Server environments, if one cluster includes multiple DRS rules with duplicate names, actions can fail with an exception. For example, duplicate-named DRS rules that control placement on storage can result in a `StoreOperationException` error.

- **Fixed Issue:**

Customer Issue 111948

Memory management for discovered Bind to Complimentary Group placement policies must be improved. In very Container large environments, Turbonomic can discover many of these policies, and memory issues can cause certain plans to fail.

- **Fixed Issue:**

Customer Issue 111948

Memory management for the history component must be improved. Under some circumstances for large environments, running nightly plans can cause the history component to run out of memory.

- **Fixed Issue:**

Customer Issue 111418,111640

For Nutanix environments, when you add a Nutanix Prism Central target the target validates but it does not discover entities.

- **Fixed Issue:**

Customer Issue 111639

For AWS environments, when you add an AWS Billing Target, then Turbonomic should discover and display all billing data. However, it does not display the associated billing families when you try to apply a Price Adjustment, unless you add the same target as an AWS Target.

- **Fixed Issue:**

Customer Issue 111418

Under some circumstances, Turbonomic fails to validate Nutanix targets.

- **Fixed Issue:**

Customer Issue 111354

When configuring a plan, the rate of resize is always the rate that is specified for the real-time market. The user interface doesn't allow you to specify a different rate of resize for a plan.

- **Fixed Issue:**

Customer Issue 110965,112012,112028

Performance for displaying groups and group-related must be improved.

Known Issues

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

For Embedded Reporting, the VM Rightsizing Report only covers On-Prem recommendations.

- **Known Issue:**

Customer Issue 112178

For public cloud environments, under some circumstances analysis can recommend an action to move a VM off of an RI and increase cost, even if that move does not increase performance.

This can happen when RIs support varying combinations of CPU and MEM size, and are available to the different instance types. For example, analysis might choose to reduce CPU, but then the available RI does not support the required MEM requirement. In that case, the action recommends moving off of the RI, which increases cost.

- **Known Issue:**

For ServiceNow environments, if an action approved in ServiceNow has then been rejected or has failed in Turbonomic, Turbonomic does not regenerate the action and send it back to ServiceNow.

24 hours after a rejection or failure, Turbonomic should check whether the action is still valid in the environment. If the action is still valid, and if you have turned on the option to retry actions in ServiceNow, then Turbonomic should send the action back to ServiceNow for another round of approval. However, it does not retry the action.

- **Known Issue:**

For ServiceNow environments, Turbonomic fails to save any automation policy that sets the Action Type to **Request Approval from ServiceNow**.

- **Known Issue:**

For AppDynamics environments, Turbonomic cannot discover databases if the target authentication uses `oAuth` for credentials.

- **Known Issue:**

When calculating reservations in the PLACE page, if you select a Merge Policy as a placement constraint, then the placement will fail.

Note that if you have created a merge policy for a given cluster, and that cluster is in the scope of your placement request, then placement automatically uses that merge policy. You should not choose the merge policy as a placement constraint.

- **Known Issue:**

For Application Component automation policies, the user interface allows you to make conflicting settings. The Action Generation setting can show incorrect values that you can choose for the policy. As a result, you cannot save the policy.

- **Known Issue:**

For ServiceNow environments, the Turbonomic user interface allows you to set orchestration for actions that the ServiceNow integration does not support. If you configure orchestration for these actions, then either the action never generates a ServiceNow CR, or the action can fail when the CR is approved.

The actions you can configure but will not generate a CR are:

- Storage Suspend
- VSan Storage Resize

Note that storage resize for a VSan is accomplished by provision/suspend of Host. You should not configure ServiceNow orchestration for VSan Storage Resize. However, Host Provision is not currently supported for ServiceNow orchestration (see next).

- Host Provision
- File Delete
- Application Component - No actions are supported

The actions you can configure but that can fail include actions that you must also configure for execution on the affected targets. These actions include:

- Host Suspend

For this action to succeed, it must be enabled in the given hypervisor, and there must be no VMs currently running on that host.

- Storage Provision

Currently Turbonomic can only execute a CR for this action on Pure and Dell Compellent storage.

- **Known Issue:**

For updates from versions earlier than 7.22.4, the update does not fully migrate policies for Application and Application Server entities.

Starting with version 7.22.4, the supply chain for applications has changed. Application and Application Server are now represented by the Application Component entity type. If you had created policies that affect these older entity types, then many of the settings will revert to their defaults. Before updating to the new version, you should check for affected policies and record the settings.

- **Known Issue:**

When you create reservations (in the Placement page), if you provide a Network constraint the reservation does not recognize that constraint. The user interface displays a notification that the reservation cannot be created. However, the platform does create the reservation, and it does not include the network constraint.

- **Known Issue:**

For vSAN environments, under certain circumstances a plan to add workloads can fail to place workloads, or it can fail to generate actions to increase storage capacity by provisioning new hosts.

- If you scope the plan to a user-created group that only provides vSAN storage, or to a discovered storage cluster group, then the plan can fail to place VMs with multiple volumes. This can occur for VMs that use conventional storage (not vSAN) along with vSAN storage.
- If you scope the plan to a vSAN host group and add workloads, the plan can fail to increase storage capacity by provisioning new hosts. For example, assume you scope the plan to a vSAN host group and add 20 VMs to the environment. In that case, you need hosts to provide compute capacity for the VMs, and you also need hosts to provide storage capacity. The plan can represent the compute provisioning correctly, but it can incorrectly fail to add more storage capacity to the vSAN.
- If the vSAN RAID type is `Raid6/FTT=2`, if you scope the plan to any vSAN groups then the plan will fail to place any of the VMs.

- **Known Issue:**

For Azure environments, when you inspect resource groups, Turbonomic does not currently show the billed costs for those resource groups.

- **Known Issue:**

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.

- **Known Issue:**

Customer Issue 112077

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

To help you visualize and interact with the REST API resources, Turbonomic includes the open source tool, SwaggerUI. For REST API developers, if you use the SwaggerUI to try out API calls, if the API call returns an error (for example, your call uses incorrect syntax), under some circumstances SwaggerUI does not display the error. The field for the returned data is empty. This appears to be a problem with the Swagger UI tool.

- **Known Issue:**

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`

- **Known Issue:**

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:**

The All Actions chart does not include pending actions for databases or database servers.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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
```

- **Known Issue:**

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}
```

- **Known Issue:**

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.

- **Known Issue:**

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

- **Known Issue:**

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 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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.

- **Known Issue:**

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`.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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.

- **Known Issue:**

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 antiaffinity rules.

- **Known Issue:**

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.