



Turbonomic 7.21.4 Release Notes

April 15, 2020

This document describes issues that are addressed in Turbonomic 7.21.4 – Release Date: April 15, 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 7.21.4

This release of Turbonomic includes the following new features:

- Application Resource Management for Microsoft Azure Environments

This release adds support for Microsoft Azure public cloud environments. To manage your Azure environment, you configure two types of target:

- Azure Service Principal Targets

Microsoft Azure is Microsoft's infrastructure platform for the public cloud. You gain access to this infrastructure through a service principal target. Service principal targets automatically discover the subscriptions to which the service principal has been granted access in the Azure portal.

- Microsoft Enterprise Agreement Targets

An Enterprise Agreement target enables Turbonomic to use custom pricing and discover reserved instances in your environment. When you configure an Enterprise Agreement target, Turbonomic uses that richer pricing information to calculate workload size and RI coverage for your Azure environment.

For more information, see "Microsoft Azure" in the *Target Guide* and "Microsoft Enterprise Agreement" in the *Target Guide*.

- Optimize Cloud Plans

This release introduces the Optimize Cloud plan. This plan identifies ways to optimize your costs by choosing the best workload tiers (most adequate compute resources), regions, accounts, or resource groups to host your

workloads. The plan also identifies workloads that can change over to RI pricing plans, and it compares your current costs to the costs you would get after executing the plan recommendations. For more information, see "Optimize Cloud Plan" in the *User Guide*.

- **Enhanced Descriptions for Delete Volume Actions**

For public cloud environments, Turbonomic can discover detached storage volumes and generate Delete Volume actions. In this way, you can use Turbonomic to make sure you do not pay for unused storage on the cloud.

To help you make decisions about these actions, with this release the details for Delete Volume Actions now include:

- How long ago the volume was last attached, in days
- The name of the VM to which the volume was last attached

- **Audit Log Improvements**

The audit log now includes entries for changes to Turbonomic policies. If a user deletes, adds, or modifies a policy for any scope in your environment, Turbonomic enters that event in the audit log.

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 environments that include SNMP and WMI targets, this release improves logging for target validation and incremental discovery, as follows:
 - On the Target Configuration page, the details for these targets now show the total number of VMs included in discovery, the number of VMs with partial discovery, and the number of VMs with failed discovery
 - For logging set to `INFO`, Turbonomic logs incremental discovery details, including the list of VMs that fail discovery, and which statistics are missing for partial failures.
- **Customer Issue 109351**

Improvement: In the Cloud view, the Locations chart now only shows locations for which there are discovered workloads.
- **Customer Issue 108596,108994**

Improvement: When you run a plan, it is possible for the plan results to include some number of unplaced workloads. This indicates that the workloads cannot be placed in the scope of your environment that the plan is configured to run against.

With this release, for any unplaced workloads, the plan results now include a description of why the given workload is unplaced. For example, an unplaced VM entry might say *Due to CPU*, to indicate that there were no available hosts with the CPU capacity that VM requires.
- **Improvement:** The audit log now includes entries for changes to Turbonomic policies. If a user deletes, adds, or modifies a policy for any scope in your environment, Turbonomic enters that event in the audit log.
- **Customer Issue 107522**

Improvement: In vCenter Server environments, it is possible to configure a VM so that it reserves all of its memory. To do this, in the vCenter user interface you enable the setting, **Reserve all guest memory (All locked)**. To support this configuration, Turbonomic respects the setting, and it will not recommend Memory Reservation resize actions for such a VM.

For example, when you have configured a VM in this way, it is possible that the VM leaves a significant amount of the reserved memory underutilized. In this case, to respect the **Reserve all guest memory (All locked)** setting, Turbonomic does not recommend a resize down action for Memory Reservation.

Fixed Issues

- **Customer Issue 110388**

Fixed: For Datadog environments, under some circumstances discovery halts before the probe discovers the full environment. This can occur when the Datadog discovery encounters incomplete data.
- **Fixed:** For offline installations, under some circumstances the Kubernetes installation script `/opt/local/bin/t8c.sh` fails to execute completely.
- **Customer Issue 110247**

Fixed: If you have updated Turbonomic from a 7.17.x version, under some circumstances plans can stop running before they have completed.
- **Fixed:** A security vulnerability was discovered in a Turbonomic component pertaining to [CVE-2018-20200](#).

- **Customer Issue 110226**

Fixed: For Hyper-V environments, under some circumstances discovery fails to discover the IPs of some VMs.

- **Customer Issue 110010**

Fixed: In very large environments, when you use the Place features to create a reservation for the global scope, under some circumstances the plan that calculates the reservation placement can fail. In this case, Turbonomic can post a `type: 500` exception to the log.

- **Customer Issue 109250**

Fixed: For vCenter Server environments, if you enable migration across vCenter instances, the migration should always keep the migrated VM on the same resource pool when applicable. For VMs managed by vCloud Director, migrations respect the resource pools. However, for environments that don't use vCloud Director the migrations don't always respect the resource pools.

- **Customer Issue 110139**

Fixed: When you configure a plan with a specific scope, the plan calculations can use providers that are outside of that configured scope.

- **Customer Issue 110125**

Fixed: In Hyper-V environments, under some circumstances when a resize down action fails for some reason, Turbonomic does not restart the VM (in its original size).

- **Customer Issue 110106**

Fixed: For Azure environments, when running a Migrate To Cloud plan, under some circumstances VMs with large disks fail to place in the "Consumption" plan for the migration. This occurs when the required storage capacity is only available on managed storage.

- **Customer Issue 110097**

Fixed: For Dynatrace environments, under some circumstances Turbonomic can discover incorrectly high heap capacity for a JVM. This can result in incorrect resize down actions.

- **Fixed:** For API development, when you get paginated data, the returned data now includes the header, `x-total-record-count` that you can use to ensure that you have retrieved all of the paginated data. However, this header is not included in all paginated data.

- **Customer Issue 110093**

Fixed: Under some circumstances, logging unnecessarily shows the error:

```
CommodityStorageprovisionedVmStBoughtCommConverter
```

This can occur when a VM has RDM disks on several datastores. This is a normal situation, so the log should not post the error.

- **Customer Issue 110063**

Fixed: For Horizon VDI environments, discovery can fail with the following exception:

`IllegalArgumentException`. This can happen when the discovered entity contains special characters in its identifier. Discovery should not generally fail when it encounters a bad identifier, and discovery should handle special characters in the identifiers.

- **Customer Issue 110056,110315**

Fixed: Under some circumstances, the historical data component experiences performance degradation related to thread pool management.

- **Fixed:** Under rare circumstances, when updating to a new version the migration of historical data can fail. This can occur because of changes between daylight savings time and standard time. Under rare circumstances some data files can get a timestamp for an hour that was skipped because of the changeover.
- **Customer Issue 110017**
Fixed: The calculation for Garbage Collection Time is incorrect, which can result in incorrect actions to manage Heap allocation. For example, if Heap utilization is low but collection time is high, then Turbonomic can recommend an increase of heap. But this example indicates correct operation for a heap that is not overutilized.
 To fix this issue, the calculation should be 100% minus the Garbage Collection percentage. In the above example, the high garbage collection time would result in a low value, and so it would not trigger a resize of the heap.
- **Customer Issue 109989**
Fixed: For Hyper-V environments, under some circumstances the execution of VM Move actions can fail to execute.
- **Customer Issue 109768,109967**
Fixed: For Datadog environments, under some circumstances Turbonomic can stop discovering Datadog metrics.
- **Fixed:** When you use the API to get information about targets, you can pass a parameter for `environment_type`. If you set this parameter to `ONPREM`, then the API request can fail.
- **Customer Issue 109668**
Fixed: For AWS environments that use Secure Token Service (STS), if you disable the `us-west-2` STS endpoint, then Turbonomic cannot validate the AWS targets. Turbonomic should not rely on an STS endpoint that it allows you to disable. Because it does not allow you to disable `us-east-1`, Turbonomic should rely on that endpoint when resolving STS authentication.
- **Customer Issue 109648**
Fixed: When you set Storage Latency charts to a two-hour time range, the charts do not show different data for Peak and Average values.
- **Customer Issue 109522**
Fixed: For VMware Horizon View targets, under some circumstances authentication does not accept valid user names.
- **Customer Issue 109473**
Fixed: For Multiple Resources charts that are scoped to a group of entities, if the group members have different resources, then it's possible that the chart will not display data for the full aggregate of resources across all members in the group.
- **Customer Issue 109486**
Fixed: When running plans, under some circumstances when the plan scenario increases resource utilization for a VM on a host, the plan can recommend suspending that host. This is inconsistent, because you should not suspend a host if you increase utilization on it.
- **Fixed:** For Google Cloud Platform environments, Turbonomic does not show cost data for discovered VMs.
- **Customer Issue 109098**
Fixed: For environments that use VMAX storage, under some circumstances Turbonomic discovery does not add all the VMAX entities to the associated disk array in the supply chain.

Known Issues

- 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.

- When you run an Optimize Cloud plan and view the RI Inventory chart, under some circumstances the *Show All* view can appear with no content. This can occur when you set the plan scope to a group.
- Customer Issue 110334,110337,110347**

Under some circumstances with very large environments, the topology-processor component can fail to start up because it is blocked by the history pod. When it is in this state, it continually tries to restart.

If this occurs, you can delete the history pod. The topology-processor should successfully restart, and the history pod should also restart. To delete the history pod, use the command:

```
kubectl delete pods history-5d595597f-kpndc
```

where

```
-5d595597f-kpndc
```

is a generated ID attached to the pod name.

For assistance, please contact Technical Support.

- 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.

- When performing an offline update of Turbonomic, under rare circumstances the upgrade process can hang on the MariaDB component. If the update process hangs, execute the following command to inspect the database component:

```
systemctl status mariadb.service
```

The database component can show messages similar to the following:

```
[Note] Event Scheduler: Stopped
```

```
[Note] Event Scheduler: Purging the queue. 13 events
```

```
[Note] InnoDB: FTS optimize thread exiting.
```

```
[Note] InnoDB: Starting shutdown...
```

```
[Note] InnoDB: Waiting for page_cleaner to finish flushing of buffer pool
```

If this occurs, execute the following commands to break out of the problem and restart the database component:

```
mysqladmin -uroot -pvmturbo shutdown
```

```
service mysql start
```

On a successful restart of the database component, the update will resume.

- 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.
- New Minimum Requirement:**

The machine that hosts the Turbonomic platform must support the SSE4.2 instruction set. Support for this instruction set was introduced at different times for different chip manufacturers:

- Intel: November 2008
- AMD: October 2011

The machine you use to host Turbonomic should be newer than these dates. On a Linux system, you can run the following command to check for this support:

```
cat /proc/cpuinfo | grep sse4.
```

For more information, see the glossary entry at <http://www.cpu-world.com/Glossary/S/SSE4.html>.

- 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 some circumstances, the writing of entity data to the historical database can fail. This can occur when the discovered entity name exceeds the character length allowed for that field in the database. When this occurs, Turbonomic logs an error that says `Data too long for column...`
- For public cloud environments with billing families, if you set the scope to a billing family, the user interface displays a count of workloads for the billing family master account. The master account should not show that it contains workloads. Instead, the billing family member accounts should be the ones to show counts of workloads.
- 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.

- 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.
- For public cloud environments that use New Relic, at a maximum Turbonomic can discover 2000 VMs that host applications managed by New Relic.
- 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.

- For AWS environments that use dedicated tenancy, under some circumstances Turbonomic can recommend changing a workload to a valid instance type for the tenant, even though the current type is already valid. This can happen when the instance type is not included in the Offer File for the tenancy.

For example, assume the `t3a` template family does not support dedicated tenancy. However, assume that the user created a `t3a` instance with dedicated tenancy in the EC2 console. In that case, Turbonomic will see this as a misconfiguration and recommend changing to a different instance type.

- For environments that use Raw Device Mapping (RDM) storage, it is possible for the user interface to show storage counts that seem inconsistent. This is because a VM can have all the potential RDM blocks as providers, while a Host has only the specifically used RDM blocks as providers.

For example, when you scope the view to the VMs on a host, the supply chain can show a higher count for storage than when you scope the view to the given host. Or when you run a plan to merge clusters, the plan results can include storage that is outside of the merged clusters.

Note that this only affects display in the user interface. Turbonomic analysis in real-time and in plans uses the correct storage relationships.

- 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.
- 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 spot instances of VMs, the Turbonomic does not discover price data for the VM instances. The user interface does not show pricing, and the pricing is not available for analysis.
- 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.

- For public cloud environments, under some circumstances you can see conflicts between the costs shown in the Cloud Cost Comparison chart, and charts for individual cloud accounts.

RI Savings for Individual Accounts

Specific RIs can provide savings for multiple accounts. Charts for an individual account show the full RI savings, which can result in exaggerated savings for that account. If you add up savings in the affected individual charts, the result is greater than the overall RI savings you would see in the Cloud Cost Comparison chart. Note that setting the Cloud Cost Comparison chart to a global scope is the best way to view the total impact of Turbonomic actions on your environment.

Disabling Resize for VM Resources

To resize VMs on the cloud, Turbonomic uses Move actions to run the VM on a different instance type. For on-prem environments it uses resize actions for specific resources. If you want to disable resizing for public cloud VMs, you should disable move actions. If you disable specific resize actions for VMs on the public cloud, the resulting display of savings can be inconsistent between the Cloud Cost Comparison chart and other savings/investment charts for the affected scope.

- For AWS environments, if a target account includes a policy with region-level control, then discovery of that target will fail with an `authorization` exception.

- **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.

- For public cloud environments, the billing data that the public cloud service makes available to Turbonomic is not complete for the current day. As a result, the RI Coverage and RI Utilization charts can show incorrect values during this time. For example, the RI Coverage chart might show less than the actual coverage because it does not have the complete RI data.

This does not have any impact on the actions that Turbonomic generates.

To address this issue, set the time range to end on the previous day for your RI Coverage and RI Utilization charts.

- 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.