Turbonomic 7.22.4 Release Notes

July 13, 2020

This document describes issues that are addressed in Turbonomic 7.22.4 – Release Date: July 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 7.22.4

Version 7.22.4

This release strongly emphasizes our application-driven approach to managing your infrastructure. With the new application-centric features in this release, it is easier for you to see the health of your environment and evaluate actions from the perspective that matters – Application Performance.

This release introduces:

- **The APPLICATION View**

  The Home Page now includes a new APPLICATION view that focuses on application performance and shows how your applications are affected by risks to the underlying infrastructure. The view includes charts for Top Business Applications, Top Business Transactions, and Top Services. You can examine underlying performance and compliance risks, and execute actions to address these risks.

  See "APPLICATION View" in the *User Guide*.
- **New Application Entity Types**

  This release introduces the following entity types to model your applications:

  - **Business Application** – The complete application as users see it. See "Supply Chain - Business Application" in the *User Guide*.

- **Business Transaction** – A user-facing capability within your Business Application. You can monitor performance as experienced by your end users in the context of Business Transactions. See "Supply Chain - Business Transaction" in the *User Guide*.
- **Service** – One or several *Application Components* (see next item) that perform specific functions for your Business Application. You can measure performance as experienced internal to the Business Application in the context of Services. See "Supply Chain - Service" in the *User Guide*.
- **Application Component** – A unit of processing within a Service that consumes resources to enable it to perform its function for the Business Application. See "Supply Chain - Application Component" in the *User Guide*.

Turbonomic automatically discovers these entities for your applications when you add the following targets:

- AppDynamics
- Dynatrace
- New Relic
- Application Insights

In the supply chain, when you hover on an entity type, you would see the percentage of risks for the underlying infrastructure. When you scope to an entity type, you can see all the actions that are valid for the infrastructure that supports that scope. You can drill down to individual entities for a more granular view.

- **Action Visibility for Applications**

  To assure application performance, Turbonomic recommends actions on the underlying *Application Components* and *nodes* (such as database servers, VMs, and containers) that provide resources to your Business Applications. The Pending Actions chart at every level of the application model shows these actions, thus providing visibility into the underlying risks that have a direct impact on your Business Applications' performance.

  > **NOTE:**
  > You can set automation and constraint policies for these underlying entities as usual.

- **Service Level Objectives (SLOs)**

  For mission-critical entities, you can set SLOs in policies and then track actual performance against those SLOs in charts. For information about setting SLOs in application policies, see:

  - "Analysis Policies: Business Applications" in the *User Guide*
  - "Analysis Policies: Business Transactions" in the *User Guide*
  - "Analysis Policies: Services" in the *User Guide*
  - "Analysis Policies: Application Components" in the *User Guide*

- **Turbonomic as a Business Application**

  The supply chain also models Turbonomic as a Business Application so you can monitor the performance of the product.

  When you scope to the Turbonomic entity, you can see the Business Transactions, Services, and Application Components that make up the product, the underlying nodes (such as the product's VM and database server), as well as performance metrics.

  For more information, see "Turbonomic as a Business Application" in the *User Guide*.

- **Data Ingestion Framework (DIF)**

  DIF enables you to define custom entities and entity metrics for your environment, and load them into the Turbonomic supply chain for analysis. DIF consists of:

  - A JSON schema that declares the structure of the data you will load into Turbonomic

- A probe component running on the Turbonomic platform that periodically requests your data to ingest any updates

We provide the JSON schema along with reference documentation on GitHub at:

https://github.com/turbonomic/turbo-go-sdk/tree/master/pkg/dataingestionframework/schema

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:** With this release, we have improved the way we handle database servers and applications (including business applications, business transactions, and services) that do not stitch into the underlying supply chain infrastructure. In that case, Turbonomic treats these entities as HYBRID, and displays them in both the ON-PREM and the CLOUD views.

  Once Turbonomic can stitch these entities into the underlying infrastructure, it then classifies the entities according to the class of the infrastructure. They will then display as either ON-PREM or CLOUD entities.
- **Improvement:** This release includes policy settings to configure Response Time SLO and Transaction SLO. You can specify an upper limit for transactions or for response time.

  You can also set whether to turn SLO on or off. When you turn it off, Turbonomic uses an auto-generated value for the SLO upper limit. SLO is turned off by default.

> **NOTE:**
> We are delivering the new SLO settings for policies as *Public Preview*. You can experiment with these settings to chart how applications meet SLO requirements for a given scope.

Fixed Issues

- **Fixed:** For very large Dynatrace environments, Turbonomic can fail to load the discovered environment data.
- **Customer Issue 111642**

  Fixed: Under rare circumstances, configured targets can repeatedly lose connection, and the targets remain in the Not Validated state.
- **Customer Issue 111609**

  Fixed: In the action details for RI Purchase actions, a chart shows the ratio of on-demand workloads to workloads covered by RI. For large environments the RATIO axis does not always display correct values.
- **Customer Issue 111502,111739**

  Fixed: For the Top Utilization charts (Top Clusters, Top Datacenters, Top Hosts, etc.), if you scope the chart to a group of groups (for example, a group of clusters), then the chart does not display data.
- **Customer Issue 111513**

  Fixed: When configuring user authentication via Active Directory (AD), it is possible to specify a specific LDAP server. If you subsequently remove the LDAP server specification, then authentication should use the entire AD domain. However, Turbonomic does not save the change when you remove the LDAP server specification from the user interface.
- **Fixed:** Plans to add host entities by template fail if you use HCI Host templates.

  When you create a plan, you can choose to add new hosts to an environment, or you can choose to replace existing hosts with a different host type. One way to specify the new hosts is to use Host templates. If you use an HCI Host template, then the plan will fail. If you use normal Host templates, the plan can succeed.
- **Customer Issue 111466**

  Fixed: For very large environments, it is possible that requests for data passed between the platform components and the user interface can exceed the set limit. This can cause the user interface to fail to display data. For example, the supply chain can be blank.

  However, configuring a greater message limit does not fix the problem.
- **Customer Issue 111361**

  Fixed: For Hyper-V environments, VMs that are configured with static memory allocation appear in Turbonomic with 100% utilization, regardless of the actual utilization.
- **Customer Issue 111611,111625**

  Fixed: For AWS multi-region environments with Autoscaling, under some circumstances Turbonomic can generate autoscaling groups and policies with duplicate names. In very large environments, this can generate log entries that exceed the supported message size.
- **Customer Issue 111405**

  Fixed: Under rare circumstances, after deleting a target that discovers policies or groups, and subsequently configuring the same target again, Turbonomic can fail to load group or policy data.

- **Customer Issue 111407**

  Fixed: After you create a custom dashboard, if you change the dashboard name or Visibility setting, then the user interface stops displaying the dashboard content.
- **Customer Issue 111316**

  Fixed: For Hyper-V or VMM environments that include SMB storage, if the storage name includes upper-case characters then in some circumstances Turbonomic does not discover latency on the affected storage.
- **Customer Issue 110995**

  Fixed: The Headroom chart is only valid for scopes of on-prem clusters. However, the user interface enables you to choose any clusters, and enables you to save the chart with an invalid cluster.
- **Customer Issue 111399**

  Fixed: In the Target Configuration user interface, it is possible to add multiple instances of the same target. You can do this by declaring the same target address, but changing the case of some of the characters. Duplicate targets can generate an invalid supply chain.
- **Customer Issue 111370**

  Fixed: 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.
- **Fixed:** For VSan environments, under rare circumstances a number of vertical scaling recommendations can generate an invalid series of events. As a result, the market component can fail with a null pointer error.
- **Customer Issue 110097**

  Fixed: For Dynatrace environments, under some circumstances Turbonomic can fail to discover heap metrics.
- **Customer Issue 111226,111307**

  Fixed: For very large environments, must improve memory management for the Market component.
- **Customer Issue 111267**

  Fixed: Under rare circumstances, Turbonomic can stop generating actions. This can occur if the Topology Processor component crashes before end of day and then restarts after the beginning of the next day, and there is a failure to populate the percentile utilization data. In that case, the affected commodities can be marked as not to be resized.
- **Customer Issue 111196**

  Fixed: For vCenter Server environments, under rare circumstances Turbonomic can falsely identify active VMDK files as wasted storage that you can safely remove. This can occur when Turbonomic discovers different modification dates for the same file, and for one of those dates the file was not active.
- **Customer Issue 111069,111755**

  Fixed: For Datadog environments, under some circumstances after validation and discovery Turbonomic does not display any data for the Datadog environment.
- **Customer Issue 111310**

  Fixed: Under rare circumstances when working with dynamic groups based on tags, two or more groups can generate duplicate keys. When this happens then Turbonomic fails to save groups with the error:

  ```
  ERROR [grpc-default-executor-107] [GroupRpcService]     : Failed to perform operation
  ```
- **Customer Issue 111085**

  Fixed: For public cloud environments, when you export the Pending Actions list to CSV, the exported data does not include values for Volume ID, Region, or Volume size

- **Customer Issue 110949**

  Fixed: Under some circumstances Turbonomic recommends resizing of Microsoft SQL instances by negligible increments.
- **Fixed:** For installations that require a unique admin account for the database component (as opposed to the account that Turbonomic uses by default), if the database connection goes down, the admin account will fail to authenticate upon reconnect attempts.
- **Fixed:** For VSan environments, the supply chain can show VMs as connected to hosts that supply VSan storage, even though the VMs do not run on those hosts. The supply chain should not display storage-providing hosts as the machines that host such VMs.

Known Issues

- When you set the scope of the Turbonomic view to a group, you can then view the automation policies that impact the given group. If you edit a policy for that group (in Settings: Policies), and then scope the view to that group again, the policy changes do not appear in the display for that group. The display should update within ten minutes, after the next round of incremental discovery. If the condition persists, log out of your session and log in again to update the display.
- **Customer Issue 105693**

  The Headroom chart for All On-prem Hosts does not agree with the Top Clusters chart.

  Turbonomic generates the All On-prem Hosts headroom data in a nightly plan. When the plan runs, this data is correct. In the course of the day, this data can become stale. To accurately track your cluster usage, you should use the Top Clusters chart.
- For 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.
- 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.
- 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.

- Under certain circumstances an Optimized Cloud plan can fail to use newly purchased RIs, even if they match the plan environment and workloads. As a result, the costs that the plan shows can be incorrect.
- For Azure environments, when you inspect resource groups, Turbonomic does not currently show the billed costs for those resource groups.
- **Customer Issue 111396**

  For cloud environments, under rare circumstances Turbonomic can recommend resizing a VM to an instance type that is older and less capable than an equally priced instance type.

  Under most circumstances, when a cloud provider offers a new instance type that is meant to replace an older type, the provider offers it at a lower cost. In at least one instance we have seen a case with identical costs for the newer and older instance types. If this occurs, and capacity and cost are equal, Turbonomic cannot ensure that it chooses the newer instance type.

  To work around this issue, you can create an Action Automation policy that excludes the older instance type.
- The user interface includes a feature to configure email and trap notifications, and the User Guide includes a description of this feature. The user interface accepts and saves your configuration, but Turbonomic does not generate any notifications.
- After you update Turbonomic from the 7.21 version family up to the 7.22 version family, when you review saved plans the plan results do not include Storage Amount data. To regenerate the Storage Amount data, run the plans again.
- For public cloud environments that include AWS and Azure, when you run the Optimize Cloud plan with a scope that includes All Providers, the RI Coverage and RI Utilization charts do not display data for AWS. To view AWS data, scope the plan to only AWS providers.
- It is possible to set the Observation Period for Percentile utilization analysis to a value that is greater than the length of data retention for historical data. For example, if you set the observation period to 90 days, that is longer than the default 60 days of data retention.

  To use a 90 day observation period for percentile analysis, be sure to increase your data retention to 90 days as well.
- **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.
- **Customer Issue 111761**

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

  Do not use the `numStorages` statistic filter. This statistic *will* be returned when other filters or no filters are used.
- The All Actions chart does not include pending actions for databases or database servers.

- For Azure environments with VMs in Scale Sets, for any VMs that are powered off the associated storage shows a utilization of zero GB. This is an accurate presentation of the data that the Azure environment returns for such a powered-off VM. However, it is likely that some of the storage capacity is currently utilized.
- **Customer Issue 110123**

  There is a memory limit for the data you can download from the All Actions chart. For example, assume you have executed many actions over time in your environment. As a result, the list of all executed actions might exceed the data limit. In that case, downloading a CSV file from the All Actions chart will fail.
- Under rare circumstances the Turbonomic platform stops responding. This occurs when `etcd.service` fails. When it does occur, you should see the following error:

  ```
  Error response from daemon: endpoint with name etcd1 already exists in network host
  ```

  To recover from this situation, restart the docker service for the Turbonomic platform. execute the command:

  ```
  sudo systemctl restart docker.service
  ```
- Under rare circumstances, when the Turbonomic platform restarts it can fail to mount the platform storage. This occurs when the heketi pod does not start up correctly. Turbonomic uses heketi and glusterfs pods for storage, and when heketi fails to start it cannot mount the storage.

  To recognize this situation, use the following command to monitor the heketi and glusterfs pods:

  ```
  kubectl get pods -A | egrep "glusterfs|heketi"
  ```

  You should see messages similar to the following:

  ```
  Warning  FailedMount  79s   kubelet, node1     MountVolume.SetUp failed for volume
   "db" : mount failed: mount failed: exit status 1
  ```

  If this occurs, delete the glusterfs pod with a command similar to the following, where `{Unique_ID}` is the ID of the glusterfs pod:

  ```
  kubectl delete pod -n default gluster-{Unique_ID}
  ```
- When you update from 7.21.0 to this version, the update process sets your **Data Retention** setting back to the default values. If you have made custom data retention settings, you should reset them after you update.
- Updates from the 7.17 version family to the 7.21 version family cancel and delete any reservations that you have set up in the Placement view. If you require these reservations, then you should configure the reservations again in the updated version of Turbonomic
- If you are performing an *Online Update* from a 7.17 version of Turbonomic, then your update can inherit old configurations for the maximum MySql connections. In large environments, or environments with many users of the same Turbonomic instance, this can result in the error, `error code [0]; Too many connections;`.

  Your Turbonomic instance should be configured for a maximum of 151 connections. You can find this configuration in the following files on the Turbonomic instance:

  - `/etc/my.cnf.d/server.cnf`
  - `/opt/turbonomic/kubernetes/etc/my.cnf`

  To correct this issue, change the configuration to allow 151 connections. The new setting should read `max_connections = 151`. After 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.

- When you update Turbonomic from the 7.17 version family to 7.21.0, you can lose any Accepted Action charts that you have included in your dashboards and views. The update changes these Accepted Action charts to All Actions charts.

  To resolve this problem, edit the All Actions charts to change them to Accepted Actions charts.
- Update Deletes Saved Plans

  Because of changes to Turbonomic plans, when you update from the 7.17 version family to the 7.21. version family, the update process deletes your saved plans.
- For very large environments that use the WMI targets, the WMI discovery can run out of memory. To address this issue, Turbonomic has improved memory handling with WMI discovery.

  As stated in the *Target Configuration Guide*, Turbonomic recommends a maximum of 500 WMI entities per WMI target. If your WMI target manages more than 500 entities, then you can see further memory issues. If you must manage more than 500 entities per WMI target, please contact Technical Support.
- **Customer Issue 108841**

  In NetApp environments, the storage controller shows 100% utilization when there are no more disks in a SPARE state that the storage controller can utilize in an aggregate. This does not indicate that the storage controller has no capacity.
- In vCenter Server environments, charts can show that a Virtual Datacenter (VDC) uses resources at more than 100% of capacity.

  The utilization metrics that vCenter returns to Turbonomic for a VDC include utilization of resources that are reserved for vCenter overhead. However, the capacity metrics that Turbonomic discovers do not account for these reserved resources. As a result, it is possible Turbonomic shows that the VDC consumes more than 100% of capacity.
- In Azure environments, a subscription can use locked storage or locked resource groups. For such subscriptions, Turbonomic discovers incomplete data. Locked resources affect Turbonomic discovery in either of these scenarios:
  - A locked resource group

    Turbonomic discovers all the entities in the resource group, but does not discover the resource group itself. For example, in the Top Accounts chart, the Resource Groups field will show no resource groups for a subscription that has a locked resource group.
  - Locked storage

    Turbonomic discovers all the entities in the resource group except the locked storage. It also discovers the resource group.
- The Turbonomic audit log tracks all communications with the platform via HTTPS. The log entries should include the IP address of the requesting client, as well as the user account. However, the log entries do not include the IP address of the originating client.
- When you specify a schedule, you must be aware of issues for Daylight Savings Time changes. For example, assume you create two schedules, one that starts at 7:00 am and lasts 12 hours, and another schedule that starts at 7:00 pm and lasts 12 hours. If you apply those schedules to policies on the same scope, then the policies can experience a scheduling conflict as Daylight Savings time changes.
- When you use the **PLACE** page to set up a reservation or a deployment, you choose the templates to represent the workload you will deploy. The templates you choose must include an **Image** specification that gives the path to the VM package, and optional placement constraints.

  Typically, you will use templates that are discovered through your hypervisor targets. Along with discovering resource capacities for the given VM, Turbonomic should also discover the Image specification for a given discovered template. However in this version, Turbonomic does not discover the Image descriptions. In addition,

discovered templates and their image specifications are read-only. For this reason, you cannot set up placement or reservations using discovered templates.

- If you run the Alleviate Pressure plan in Turbonomic 7, and then compare it to the same plan and scope in a 6.1 release, then the display of instances in the supply chain are not identical for both versions.
- Ring charts that show the utilization of different resources show a yellow segment whenever the Reserved Capacity for the resource is zero. For some resources there is no concept of reserved capacity, yet the ring chart still shows a yellow segment.
- For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.
- **Customer Issue 99189,99805**

  In vCenter environments, you might see unusually high storage latency values or excessive recommendations to provision new storage. There is a known problem with the storage latency values that vCenter Server versions 6.5.u1x and earlier return via the API. These versions can return unusually high storage latency values.

  Turbonomic considers storage latency when calculating whether to move a VM to existing storage, or whether to provision new storage. Because of this known problem, Turbonomic can incorrectly recommend provisioning storage when moves are appropriate.

  If you encounter this problem, then you should create a policy that disables storage moves for VMs that are managed by vCenter Server versions 6.5.u1x and earlier. To create this policy:

  - Create a VM group that contains all the affected VMs. Note that Turbonomic automatically creates a group named `VMs_vCenter` that you might be able to use.
  - Create a new VM automation policy. This policy will disable storage move actions.
  - Set the group that you created to be the policy scope.
  - Under **Action Automation** add the `Storage Move` action and set it to `Disabled`.
- In cases where actions recommend that you suspend hosts, the Optimal Improvements chart should indicate no utilization on the hosts to be suspended. Under some circumstances, the chart can show utilization on these hosts. The result is incorrectly low values for utilization on the other hosts in the current scope.
- Turbonomic generates special average or max utilization templates that it uses when calculating cluster headroom. You should not edit these templates, because Turbonomic will overwrite your changes the next time it generates the templates. However, the Template Catalog presents these templates as editable.
- After you run a plan, the user interface enables you to make changes to the plan configuration and then run the plan again. If you do this, the plan results will be inconsistent. If you want to run a plan with a different configuration, you should start a new plan.
- You should never use duplicate names for groups of the same entity type. However, the user interface does not validate group names to keep you from creating a duplicate name.
- For VMs running on Hyper-V, if you set a VCPU limit (limit VCPU to less than 100%), then the VCPU utilization data that VM returns to Turbonomic is not correct. As a result, Turbonomic will not recommend that you increase the VCPU limit.
- For AWS environments, under very rare circumstances you can have RIs on payment plans that do not resolve to 1-year or 3-year terms. In this case, AWS does not return pricing data for those RIs. Turbonomic does not include such RIs in its calculations of RI utilization or RI cost.
- For vCenter Server environments, Turbonomic does not recognize DRS rules for VM restart dependencies that are based on `ClusterDependencyRule`. You might be able to achieve a similar effect by expressing dependencies l via `ClusterVmHostRule` or cluster affinity or antiaffinity rules.
- **Customer Issue 109389**

In vCenter Server environments that have Instant Clone VMs, under some circumstances Turbonomic cannot move these VMs to other hosts in the cluster, even though you can manually migrate them via the vCenter Server user interface.