



Turbonomic 7.22.10 Release Notes

January 26, 2021

This document describes issues that are addressed in Turbonomic 7.22.10 – Release Date: January 26, 2021. Please see the Turbonomic 8 documentation for earlier versions of the Release Notes:

<https://docs.turbonomic.com/>

NOTE:

These release notes are for update to Turbonomic 7.22.10, from version 7.22.0 or later. For updates from the 7.21.x family or earlier, please contact your Technical Support representative.

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 7.22.10

Version 7.22.10

- **Cloud Resource Management**

This release introduces key features and enhancements to help optimize your cloud workloads' performance and reduce your cloud expenditure.

- **Volume Scaling**

Turbonomic can now recommend scaling your AWS and Azure volumes. It uses percentile calculations to drive more accurate scaling actions, and indicates whether these actions are non-disruptive or reversible. With these capabilities, you can execute scaling actions with confidence.

For details, see "Supply Chain - Volume" in the *User Guide*.

- **Azure SQL Database Scaling**

This release introduces support for Azure SQL Databases that use the DTU (Database Transaction Unit) purchase model. Turbonomic automatically discovers these databases when you add Azure targets, and then presents them as *Database* entities in the supply chain.

Turbonomic monitors DTU and storage utilization for the databases, and uses percentile calculations to drive more accurate scaling actions.

For details, see "Supply Chain - Database" in the *User Guide*.

- **VM Scaling Enhancement**

You can now direct Turbonomic to only execute cloud VM scaling actions that improve performance or reduce costs. When you create a policy for the cloud VMs, choose the **Scale for Performance** or **Scale for Savings** option, and then set the action mode to *Automated* or *Manual*. If you set the action mode to *Manual* and then examine the pending actions for those VMs, only actions that satisfy the policy are allowed to execute. All other actions are read-only.

NOTE:

The scaling option **Cloud Compute Scale** has been renamed **Scale All**. This enables all scaling actions, including those that result in efficiency improvements and increased costs. When policy conflicts arise, **Scale All** overrides the other two scaling options in most cases. For more information, see "Relationship Between Scoped and Default Policies" in the *User Guide*.

- **Migrate to Cloud Plans**

This release introduces Migrate to Cloud plans to help accelerate and optimize your cloud migration. These plans simulate migration of on-prem VMs to the cloud, or migration of VMs from one cloud provider to another.

For details, see "Migrate to Cloud Plan" in the *User Guide*.

- **Separate Billing Targets for AWS Environments**

With this release Turbonomic introduces a separate target to discover AWS billing information. This data can be used when discovering costs for RIs and workload instances.

The separate billing target supports the practice of purchasing RIs in a master account, and running workload instances in sub accounts. Turbonomic can attach to the master account with fewer privileges than are required for the sub accounts that manage workloads.

- **Consistent Resizing**

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

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

For details, see "Analysis Policies: VMs > Enable Consistent Resizing" in the *User Guide*

- **Enhanced IOPS-aware Scaling for Azure VMs**

Turbonomic now considers IOPS utilization when making scaling decisions for Azure VMs. To measure utilization, Turbonomic takes into account a variety of attributes, such as per-disk IOPS utilization, whole VM IOPS utilization, cache settings, and IOPS capacity for the VMs. It also respects IOPS utilization and aggressiveness constraints that you set in VM policies.

For details, see "Enhanced IOPS-aware Scaling for Azure VMs" in the *User Guide*.

An IOPS-driven scaling action shows *IOPS Congestion* as the reason for the action. You can expand the action and see the *VM IOPS Percentile* section to see utilization details.

- **Improved Scaling Constraint for Cloud Workloads**

This release introduces the **Cloud Instance Types** scaling constraint in policies for VMs, databases, and database servers. This is an improvement on the **Excluded Cloud Tiers** setting in previous versions. With this change, you now select instance types for the scoped workloads and see their resource allocations right away.

- **Improvements for Potential Savings or Investments Charts**

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

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

- **Support for JVM Applications**

Turbonomic now supports JVM Applications as targets, and can recommend resizing heap and thread pool capacity. For details, see "JVM Application" in the *Target Configuration Guide*.

- **Container Management**

The following key features are now available to help you manage your container infrastructure more efficiently.

NOTE:

Deploy the latest Kubeturbo image to take advantage of these features. For information about Kubeturbo releases, go to:

<https://github.com/turbonomic/kubeturbo/wiki/Server-Versions-and-Kubeturbo-Tag-Mappings>

- **New Container Model**

This release introduces a new container model for Kubernetes environments to better reflect resource utilization and constraints, and drive more accurate actions.

NOTE:

To take advantage of the features that are part of the new container model, you must deploy a new Kubeturbo image. For information about the latest Kubeturbo releases, go to <https://github.com/turbonomic/kubeturbo/wiki/Server-Versions-and-Kubeturbo-Tag-Mappings>.

Turbonomic now discovers the following entity types. The default views for these and all other container entity types include relevant charts to help you assess the health of your container infrastructure, identify risks, and execute actions.

- **Container Cluster**

This entity type represents Kubernetes clusters in the supply chain. This allows Turbonomic to fully link your container infrastructure with the underlying nodes, and then present all actions on containers and nodes in a single view. This gives you full visibility into the actions that impact the health of your container environment.

For details, see "Supply Chain - Container Cluster" in the *User Guide*.

- **Namespace** – A logical pool of resources that manages workloads based on specific requirements or business needs. Turbonomic treats quotas defined in namespaces as constraints when making sizing decisions for containers.

For details, see "Supply Chain - Namespace" in the *User Guide*.

- **Container Spec** – A shared definition for all ephemeral container replicas. This persistent entity retains the historical utilization data of containers, which Turbonomic leverages to make container sizing decisions.

For details, see "Supply Chain - Container Spec" in the *User Guide*.

- **Workload Controller** – A Kubernetes controller that watches the state of your pods and then requests changes where needed. You can execute container resize actions when you set the scope to a Workload Controller.

For details, see "Supply Chain - Workload Controller" in the *User Guide*.

- **New Container Dashboard**

This release includes a new predefined **Container Platform Dashboard**. Use the charts in the dashboard to assess the health of your Container Clusters, Namespaces, and Services.

- **Container Resizing Improvements**

- Resizes in Workload Controllers:

Resize actions no longer display when you scope to Containers, given their ephemeral nature. Set the scope to Container Specs to view the actions. You can execute actions via Workload Controllers as usual.

To automate actions, set the action mode for Resize actions as follows:

- Container Spec policies: Manual (default)
- Workload Controller policies: Automatic

For details about automating resize actions, see "Container Actions" in the *User Guide*.

- **Merged Actions**

Executing several container resize actions can be very disruptive since pods need to restart with each resize. To minimize disruptions, Turbonomic consolidates resize actions into one merged action for replicas of the container scale group(s) related to a single Workload Controller. When a merged action 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 a merged action, set the scope to the Workload Controller for the containers or Container Spec being resized.

- **Container Policies**

Use automation policy settings to fine tune actions for Container and Container Spec entities.

- Millicores

In policies and actions, resize values for vCPU limits and requests are now expressed in millicores.

- Container

For containers, use increment constants and rate of resize to control VCPU and VMEM resizes. You can also turn on consistent resizing for container groups as a way to adhere to Container Spec settings.

For details, see "Analysis Policies: Containers" in the *User Guide*.

- Container Spec

Resize constraints previously set in Container policies are now set in Container Spec policies. Constraints include increment constants and rate of resize.

NOTE:

Turbonomic does not migrate these constraints. Record them before updating to this version, and then add them to your Container Spec policies after updating.

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

For details, see "Analysis Policies: Container Spec" in the *User Guide*.

You can also create Container Spec policies to enable tuned scaling for vCPU/vMem limits and requests. With tuned scaling, Turbonomic automates resizes if the resize values fall within a normal range, and posts more conservative actions when resize values fall outside the range.

For details, see "Tuned Scaling for Container Specs" in the *User Guide*.

- **Container Utilization Charts**

Turbonomic uses percentile calculations to measure container resource utilization more accurately, and drive scaling actions that improve overall utilization. When you examine the details for a pending scaling action on a Container Spec, which persists historical utilization data for containers, you will see charts that highlight resource *utilization percentiles* for a given observation period, and the projected percentiles after you execute the action. The charts also plot *daily average utilization* for your reference. If you have previously executed scaling actions on the associated Workload Controller, the charts also show the resulting improvements in daily average utilization. Put together, these charts allow you to easily recognize utilization trends that drive Turbonomic's scaling recommendations.

- **Support for Oracle Database Servers**

Turbonomic now supports Oracle database servers as targets, and can recommend resizing their connection and transaction log capacity. It can also recommend resizing DBMem while taking cache hit rate into account.

For details, see "Oracle Database Server" in the *Target Configuration Guide*.

- **VMware vCenter Support**

Turbonomic has certified VMware vCenter 7.0 U1 as a hypervisor target.

- **New Application Model**

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.

- **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. 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. You can monitor overall performance, make resourcing decisions, and set policies in the context of your Business Applications.

For details, 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.

For details, 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.

For details, 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.
- For details, 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. You can drill down to individual entities for a more granular view.

- **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*

- **Action Visibility**

To assure application performance, Turbonomic recommends actions on the *Application Components* and *nodes* (such as database servers, VMs, and containers) that provide resources to your 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 application performance.

- **User-defined Application Topology**

You can create your own *Business Applications*, *Business Transactions*, and *Services* without the need to ingest additional application data into the platform. This is especially useful in cases where there are gaps in the application stack shown in the Turbonomic supply chain. For example, in the absence of an application monitoring target such as AppDynamics or Dynatrace, you will not see Business Applications in your supply chain. This feature addresses those gaps.

When you create a new application entity, you identify interrelated application entities and nodes in your existing environment for which you want to measure performance, so Turbonomic can link them in a supply chain and represent them as a unified group. You can monitor overall performance for the group in the context of the new application entity.

For details, see "Application Topology" 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>

- **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 details, see "as a Business Application" in the *User Guide*.

- **Application and Database 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

- **On-Prem Resource Management**

- **Integration with VMware Horizon VDI**

To support virtual desktop infrastructures, Turbonomic has added new entities to the Supply Chain:

- **Business User** – A user account that is entitled to launch one or more active VDI sessions.
For details, see "Supply Chain - Business User" in the *User Guide*.
- **Desktop Pool** – A collection of desktops that support users.
For details, see "Supply Chain - Desktop Pool" in the *User Guide*.
- **View Pod** – A logical grouping of desktop pools, that provides a capacity for active sessions.
For details, see "Supply Chain - View Pod" in the *User Guide*.

- **Horizon VDI Used Sessions**

Turbonomic tracks two states of used sessions for your Horizon VDI desktop pools - *Active* and *Disconnected* (non-terminated). To view relevant data, scope to a Desktop Pool or View Pod and see the following charts:

- **(*New) Total Sessions** chart – Plots used sessions (Active and Disconnected) over time
- **Active Sessions** chart – Plots Active sessions over time
- **Capacity and Usage** chart – Shows current Total and Active sessions

This release also introduces a new *Desktop Pool Active Sessions* search filter for VMs. Setting the value to TRUE returns a list of VMs on your desktop pools with active sessions, while a value of FALSE returns all VMs with used sessions.

- **Improved Metrics Discovery for vCenter Server Environments**

For vCenter Server environments, this release introduces an option to collect advanced Guest memory metrics.

When you specify a target for vCenter Server version 6.0 or later, you can turn on the **Enable Guest Metrics** option. Discovery will then access advanced metric data that the target monitors for its workloads. This can increase the accuracy of the VMEM data that Turbonomic uses for analysis of VMs.

Note that VMware Tools must be installed and running on the affected VMs. The **Hypervisor VMEM for Resize** option must also be turned on for the VMs discovered by the affected target. You make this setting in the Actions section of the automation policy for the given scope of VMs. This is turned on by default.

For details, see "Hypervisor VMEM for Resize" in the *User Guide*.

- **Improved vSAN Support for Hyperconverged Environments**

For environments that use hyperconverged infrastructure to provide storage on vSAN, Turbonomic can represent the storage that is provided by a cluster of HCI hosts as a single storage entity. Analysis can generate actions to scale Storage Amount, Storage Provisioned, and Storage Access in the vSAN. You can use HCI Host templates to run plans that add capacity to your vSAN.

For details, see "vSAN Storage" in the *User Guide*.

- **Nutanix Support**

This release introduces support for Nutanix hyperconverged infrastructures. This includes management of Nutanix fabrics, where the supply chain treats a Nutanix Storage Pool as a disk array. Turbonomic recognizes Nutanix storage tiers when calculating placement of VMs and VStorage, and it can recommend actions to scale flash capacity up or down by adding more hosts to the cluster, or more flash drives to the hosts.

- **Reservation Improvements**

This release improves the performance when you create reservations. Turbonomic calculates placement more quickly, and it catches errors such as duplicate names for entities.

For API developers, the call to create a reservation is the same, except you do not pass the `apiCallBlock` parameter at creation time to make the reservation call synchronous. Instead, you create the reservation, and then execute a call to `GET` the reservation by ID, and pass the parameter `callBlocking=true`.

- **Entity Placement Constraints**

When you drill down to a single entity, you can see details about the entity's relationships in the supply chain. This shows you which entities provide resources to the entity, and which other providers can be available to it. For each listed entity, you can open a **Constraints** fly-out that gives more details about limits on the current entity's supply chain relationships.

For details, see "Entity Placement Constraints" in the *User Guide*.

- **Improved Actions for Hosts in Maintenance Mode**

This release improves performance of action management when a host goes in or out of maintenance mode. Now Turbonomic more quickly propagates the change of host state throughout the system, and clears the impacted actions from the Pending Actions list.

- **Shared-Nothing VM Migrations**

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

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

- **Action Automation and Orchestration**

- **Improved Workflow for Creating Automation Policies**

This release introduces a more intuitive workflow for creating automation policies. It is now easier to set action *automation modes*, *schedules*, and *orchestrators* (such as ServiceNow).

You can now assign a schedule to action execution, as a way to defer execution to a non-critical time window. For example, if a workload experiences memory bottlenecks during the week, you can defer the necessary resize to the weekend. Even if the workload has minimal utilization over the weekend, Turbonomic can recognize the need to resize, and will execute the action.

For details, see "Creating Scoped Automation Policies" in the *User Guide*.

- **Native Integration with ServiceNow**

This release introduces *Turbonomic Actions*. This is a ServiceNow application that you install via an update set. Turbonomic Actions integrates Turbonomic with your ServiceNow account to log the actions Turbonomic generates, and to defer these actions to your ServiceNow workflow for approval. Upon approval, Turbonomic executes the actions at the scheduled time, and logs the result in ServiceNow.

Turbonomic Actions uses the **Execution Schedule** that you assign to actions in Turbonomic automation policies.

For information about setting up policies for these actions, see "Creating Scoped Automation Policies" in the *User Guide*.

To get the Turbonomic Actions application, contact your sales representative.

- **Charts and Reports**

- **VCPU and VMem Utilization Charts for VMs**

Turbonomic uses percentile calculations to measure VCPU and VMem utilization more accurately, and drive scaling actions that improve overall utilization and reduce cost for cloud VMs. When you examine the details for a pending scaling action on a VM or an optimization action in Migrate to Cloud plans, you will now see charts that highlight VCPU and VMem *utilization percentiles* for a given observation period, and the projected percentiles after you execute the action. The charts also plot *daily average utilization* for your reference. If you have previously executed scaling actions on the VM, you can see the resulting improvements in daily average utilization. Put together, these charts allow you to easily recognize utilization trends that drive Turbonomic's scaling recommendations.

For on-prem VMs, you will see either a VCPU or VMem chart, depending on the commodity that needs to scale. For cloud VMs, both charts display.

These charts also appear when you scope to a given VM (on-prem or cloud) and view the Details page.

- **Embedded Reports**

This release introduces Turbonomic Embedded Reports. Use this feature to understand application resource management trends, and to share insights with stakeholders via reports and dashboards.

Embedded reporting runs in its own component, as part of the Turbonomic platform. This architecture enhances performance and reduces storage requirements. When you enable Embedded Reports, you can navigate to a set of pre-built dashboards that chart multiple environment details. Dashboards and charts are powered by the Grafana observability platform. With Grafana, it's easy to navigate the existing dashboards, and to make your own charts and dashboards with no coding required.

We have added presentations for:

- Wasted Storage Files
- VM Rightsizing Recommendations
- VM Utilization Day over Day By Hour
- VM Utilization Week over Week By Day

For information about enabling Embedded Reports, see the *Turbonomic Installation Guide*.

- **User Interface Administration**

- **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.

- **Feedback and Diagnostics**

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

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

For details, see "Feedback and Diagnostics" in the *User Guide*.

- **Operational Observer Role**

Users with the new Operational Observer role can view the environment, including the Home Page, Dashboards, Groups, and Policies. They can also use Search to set a scope to the session.

- **SSO Authentication Improvement**

For deployments that use Single Sign On (SSO), Turbonomic can now authenticate a user when SAML returns more than one user group for the given credentials. In the case of multiple groups, Turbonomic authenticates the user as a member of the group with the least privileges.

- **Platform Migration Tool**

This release introduces tooling to help you migrate from the 6.4 family of Turbonomic to the 8.0 family.

The 8.0 family of Turbonomic is a new platform to manage application performance to the same standards you are used to with the 6.4 version family. It introduces a component-based architecture that can manage larger environments, as well as enhancements to the supply chain and user interface that emphasize the management and performance of your applications.

Because of these changes, you must *migrate* to the 8.0 version family – You cannot perform a simple update like you would to update to a later version of the same family. With this release we introduce scripted tooling that helps you transfer your target, group, policy, and other data from one 6.4 installation, into a new 8.0 installation.

For details, see the *Turbonomic Installation Guide*.

Configuration Requirements

For this release of Turbonomic, you should satisfy the following configuration requirements.

SQL Modes for External Databases

If you deploy Turbonomic to work with an external database instead of the included historical 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, NO_ZERO_IN_DATE, or NO_ZERO_DATE.

Multi-tenant External DBs and Updates to 7.22.10 or Later

NOTE:

Azure database services are multi-tenant. If you deployed an external database on Azure, this configuration requirement applies to you.

If you are updating from Turbonomic 7.22.9 or earlier to version 7.22.10 or later, *and* you have deployed an external MariaDB or MySQL database service that is multi-tenant, then you must update the database service before you update your Turbonomic instance. This modification includes:

- Manually creating a database for the `repository` component

This includes creating the database in your DB instance, creating a user to access the database, and granting access to the user.
- Manually adding the repository database to your `cr.yaml` file

For deployments that manually add the component databases (a requirement for multi-tenant DBs), the `cr.yaml` file declares entries for each component databases. Each entry names the component, and gives the user and password that the component can use to access that database. You must add a new entry for the `repository` component.

NOTE:

If you have already updated to the latest version of Turbonomic, and you did not perform the steps to update your multi-tenant DB service, please contact your support representative.

To create the new database:

1. Connect to your multi-tenant DB using a global account.

The account must have privileges to create databases and users. If you have specified `dbRootUsername` in the `cr.yaml` file, you can use that account.
2. Create the database that the `repository` pod requires:


```
create database repository;
```
3. Create the account the pod will use to access the database:


```
create user 'repository'@'%' identified by 'vmturbo';
```

NOTE:

The value `vmturbo` is the default password that uses for all component database accounts. If you have manually created accounts with different credentials, you can do so for this database as well.

4. Set the user account privileges for the `repository` pod's account:


```
GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, REFERENCES, INDEX, ALTER, CREATE TEMPORARY TABLES, LOCK TABLES, EXECUTE, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, EVENT, TRIGGER ON repository.* TO 'repository'@'%';
```
5. Flush privileges to make the privileges take effect:


```
flush privileges;
```

Now that the repository database is created in your multi-tenant DB service, you must declare the access in the Turbonomic `cr.yaml` resource.

1. Open the `.cr` file for editing. The location of the file depends on the type of Turbonomic installation you are configuring.

OVA installation of Turbonomic:	Turbonomic on a Kubernetes node or node cluster:
<p>Open a SSH terminal session on your Turbonomic instance</p> <p>Log in with the System Administrator that you set up when you installed Turbonomic:</p> <ul style="list-style-type: none"> • Username: <code>turbo</code> • Password: <code>[your_private_password]</code> <p>Then edit the file:</p> <pre>/opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml</pre>	<p>Open the following file for editing:</p> <pre>deploy/crds/charts_v1alpha1_xl_cr.yaml</pre>

2. Add the credentials for the repository component database.

Add the entry to the `properties:` section of the `cr.yaml` file, where `vmturbo` is the password that you assigned to that user account, and `yourDB` is the qualified name of your multi-tenant DB partition::

```
repository:
  repositoryDbUsername: repository@yourDB
  repositoryDbPassword: vmturbo
```

The resulting `cr.yaml` file should be similar to:

```
properties:
  global:
    enableSecureDBConnection: true
    sqlDialect: MYSQL
    dbRootPassword: yourPassword
    dbRootUsername: xladmin@yourDB
    #dbUserPassword:
    #dbUsername:
  action-orchestrator:
    actionDbUsername: action@yourDB
    actionDbPassword: vmturbo
  auth:
    authDbUsername: auth@yourDB
    authDbPassword: vmturbo
  clustermgr:
    clustermgrDbUsername: clustermgr@yourDB
    clustermgrDbPassword: vmturbo
  cost:
    costDbUsername: cost@yourDB
    costDbPassword: vmturbo
  group:
    groupComponentDbUsername: group_component@yourDB
    groupComponentDbPassword: vmturbo
  history:
```

```

    historyDbUsername: history@yourDB
    historyDbPassword: vmturbo
  plan-orchestrator:
    planDbUsername: plan@yourDB
    planDbPassword: vmturbo
  topology-processor:
    topologyProcessorDbUsername: topology_processor@yourDB
    topologyProcessorDbPassword: vmturbo
  repository:
    repositoryDbUsername: repository@yourDB
    repositoryDbPassword: vmturbo

```

After you have done this, you can update to the latest version of Turbonomic and then apply the cr.yaml file.

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:**

Customer Issue 8632

You can now customize the format of the Turbonomic syslog output.

To customize the format, edit the cr.yaml file to specify the token order of log outputs, then use kubectl to apply the changes. For example, add the following to the cr.yaml file:

```

rsyslog:
  env:
    - name: AUDITFORMAT
      value: "<%pri%>%timestamp:::date-rfc3339% %hostname% %syslogtag%msg:::sp-if-no-1st-sp%msg:::drop-last-1f%"

```

where:

- `<%pri%>` is the log index
 - `%timestamp:::date-rtc3339%` is the timestamp and date format
 - `%hostname%` identifies the component generating the log message
 - `%syslogtag%` identifies the log message type
 - `%msg:::sp-if-no-1st-sp%` prints the message, and adds a leading space if none is present
 - `%msg:::drop-last-1f%` drops the trailing linefeed from the message
- **Improvement:**
To simplify management of Operator during Turbonomic updates, a number of properties have been removed from Operator and are now specified in component configuration files.
 - **Improvement:**
For Container entities, Turbonomic now represents CPU resources in units of `milicore`, rather than megahertz.
 - **Improvement:**
Customer Issue 113217
Starting with this release, Turbonomic now supports running the platform in an OpenShift environment that uses Security Context Constraints (SCC) to limit the valid user IDs for the platform components. The platform components now run with the assigned User IDs.
 - **Improvement:**
To simplify management of Operator during Turbonomic updates, a number of properties have been removed from Operator and are now specified in component configuration files.
 - **Improvement:**
With this release, when you configure a Migrate to Cloud plan, you can now specify which RIs the plan can use, from among the RIs that are available in the plan's scope.
 - **Improvement:**
Customer Issue 112877
For the Top Database Serves charts, you can now choose the Data Type for the chart. The choices are:
 - Health: The chart displays the health of the top utilized entities in the chart's scope.
 - Utilization: The chart displays the utilization of specific resources for the top utilized entities. You can add commodities to the chart to track their utilization.
 - **Improvement:**
Customer Issue 112814
You can now declare tags as storage annotations in the Turbonomic operator as a global entry:


```
global:
  storageAnnotations:
    turbo: tag
```
 - **Improvement:**
For Container environments, this release introduces filtering to create groups of namespaces or entities based on the labels and annotations applied to entities in the environment. With this filter, you can use namespaces to group workloads related to a given application.
 - **Improvement:**

Customer Issue 113008

For Azure environments, starting with this release Turbonomic discovers updated RI cost data via the Azure cost catalog API. Previous releases used a static CSV file that was published by Microsoft.

- **Improvement:**

For Migrate to Cloud Plans, this release now includes the RI Purchase Profile wizard in the plan configuration workflow.

- **Improvement:**

Customer Issue 111621

Starting with this release, you can specify Rate of Resize in scoped automation policies.

- **Improvement:**

Customer Issue 111367,112917

In Hyper-V and vCenter Server environments, under some circumstances a host in a cluster can have multiple D: drives. In previous versions, discovery would fail for a cluster with such a host. With this version, Turbonomic now supports discovery of hosts with multiple D: drives.

- **Improvement:**

For Kubernetes environments, when you set the scope to a Namespace entity, the charts now show the capacity and utilization of resources. This includes:

- Virtual CPU
- Virtual Memory
- Virtual CPU Request
- Virtual Memory Request

- **Improvement:**

When using **Show All** from the Savings and Investments charts, the table shows actions and now includes the ability to execute actions from that table.

- **Improvement:**

Customer Issue 110416

Starting with this release, you can now create placement policies for storage volumes and storage tiers. For example, you can now create a policy that will not place specific volumes on a given storage tier.

- **Improvement:**

The new settings `percentileAggressivenessDatabase` and `maxObservationPeriodDatabase` were added to the `marketsettingsmanager` to support Azure Single Databases

- **Improvement:**

Customer Issue 105173,108179,112638

You can now direct Turbonomic to only execute cloud VM scaling actions that improve performance or reduce costs. When you create a policy for the cloud VMs, choose the **Scale for Performance** or **Scale for Savings** option.

- **Improvement:**

Customer Issue 108292

For Oracle environments, this release introduces support for environments that deploy Oracle databases to Multiple Homes.

Fixed Issues

- **Fixed Issue:**
For Kubernetes environments, under some circumstances discovery can halt for a given cluster. This can happen when multiple pods in a node attach to the same Persistent Volume.
- **Fixed Issue:**
Customer Issue 113564,113632,113660
When you download CSV data for actions, the "From" and "To" fields are blank.
- **Fixed Issue:**
Customer Issue 113476
Under rare circumstances, updates to the historical database can stop unexpectedly. This appears in the user interface as blank intervals in the charts.
- **Fixed Issue:**
After you enable Volume Scaling actions, under some circumstances analysis can halt actions generation.
- **Fixed Issue:**
Customer Issue 113523
For AWS environments, for some RI instances analysis fails to parse the format of the RI data, and discovery for that target fails.
- **Fixed Issue:**
For a given schedule, when you click to see the list of policies that the schedule is used in, the user interface does not display any data.
- **Fixed Issue:**
Customer Issue 113377
When you create an application topology via User Defined Topology, under some circumstances the supply chain can fail to stitch the entities together.
- **Fixed Issue:**
When you scope the Workload Cost Breakdown chart to multiple types of workloads, the chart does not show reliable data. For example, this can happen if the chart is scoped to VMs and Databases. The faulty data display does not affect Turbonomic analysis.
- **Fixed Issue:**
Customer Issue 113418
For Virtual Machine Manager (VMM) environments, under rare circumstances discovery can fail when the VMM environment reports a null host cluster.
- **Fixed Issue:**
Customer Issue 112459
When you export the Top Accounts chart to CSV data, any data field that shows N/A in the chart shows 0 in the CSV data.
- **Fixed Issue:**
Customer Issue 113284

When you create or edit a chart that offers a choice of commodities, for scoped users the commodity choice is greyed out.

- **Fixed Issue:**

For Migrate to Cloud plans, if you plan a migration to an Azure environment, under some circumstances the plan results can fail to show RI costs.

- **Fixed Issue:**

Customer Issue 113401

For environments with baremetal hosts, validation of the SNMP target can fail.

- **Fixed Issue:**

Customer Issue 113379,113852

When running plans that ignore constraints, if the plan scope includes dynamic group, and the regex for that group is invalid, then the plan fails to complete. The plan should continue to execute.

- **Fixed Issue:**

Customer Issue 113388

For OpenShift environments that use DeploymentConfig to run applications, when Turbonomic executes Merged Actions (compound actions), the actions can fail.

- **Fixed Issue:**

Customer Issue 113365,114121

When you change the Default Policies to disable all actions, the setting does not take effect.

- **Fixed Issue:**

When you run a Migrate to Cloud plan with more than 500 entities to migrate, the results can incorrectly show unplaced VMs.

This can occur because the user interface requests results data for 500 entities. In a large plan, the excess entities can appear as unplaced. When you export the Cloud Cost Comparison data to CSV, that data includes the complete set of entities, showing their placement correctly.

- **Fixed Issue:**

Customer Issue

If you assign a schedule to one of the default Automation Policies, Turbonomic will respect that schedule and limit action execution accordingly. However, if you restart your Turbonomic instance, then it removes all the schedules from the default policies.

For example, assume you scheduled automation for certain VM actions in the default policy. After you restart the instance, those actions will still be automated, but they will not have a schedule. The result is that the actions are automated at all times.

After a restart, review your default policies and reapply schedules to any policies that need them.

- **Fixed Issue:**

Customer Issue 113314

For Active Directory user accounts, after updating to Turbonomic version 8.0.1, there are cases where members of AD groups cannot log in.

This can occur for members of groups where the LDAP distinguished name (DN) contains certain special characters.

- **Fixed Issue:**

Customer Issue 113294

For UCS environments, the reporting of IOModule throughput is excessively high.

- **Fixed Issue:**

Customer Issue 113196

When you configure an AppDynamic target, the user interface does not validate the fields. As a result, if you make invalid settings the target can fail to validate. For example, the target configuration should not accept an empty **Port** field.

- **Fixed Issue:**

Customer Issue 113236

For AppDynamics environments, if you use special characters in your password the target can fail to validate.

- **Fixed Issue:**

Customer Issue 113231

For Microsoft Azure environments that use Availability Sets (AS), under some circumstances resize actions can fail to execute.

According to Microsoft, if the new instance type for a resize action is not supported in the current cluster, then all the VMs in the AS must be deallocated before executing the resize. This is necessary so that Azure can identify a different compute cluster supports the instance type. If any VMs in the AS are still running, then the resize action fails.

- **Fixed Issue:**

Customer Issue 113051,113100

When using the Top Utilized charts, if you scope the chart to a dynamic group, then it does not display data correctly.

- **Fixed Issue:**

For Hyper-V and vCenter Server environments, under rare circumstances Turbonomic can experience a java deadlock.

- **Fixed Issue:**

Customer Issue 113176,113574

For HPE OneView environments, discovery can fail with `AUTHN_SESSION_CLIENT_LIMIT_CROSSED` errors.

- **Fixed Issue:**

Customer Issue 113233

For AWS environments, under rare circumstances discovery can fail and the user interface does not display any topology.

This can happen if discovery parses incorrect IDs for master and member accounts in a billing family. Discovery must query AWS for more reliable sources of this data.

- **Fixed Issue:**

For this release, when you scope the view to entities related to container environments, the Details view shows more relevant data. If you scope to Container, Pod, ContainerSpec, or WorkloadController entities, the single resource charts in the Details view now show:

- The allocated resource capacity
- Summary informaton

- Display as a Band chart

- **Fixed Issue:**

Customer Issue 113205,113434,113521

Migrate to Cloud plans have a limit to the number of VMs they can migrate. A plan to migrate thousands of VMs can fail with a `StatusRuntimeException` error.

- **Fixed Issue:**

When you run an Optimize Cloud Plan, under some circumstances the results in the OPTIMIZED column of the Cloud Cost Comparison chart can show zero entities for different counts of workloads. If this happens, try navigating away from the plan results, and back again.

- **Fixed Issue:**

For scaling actions in cloud environments, under some circumstances the action details do not indicate the commodity shortage that caused analysis to recommend the action.

- **Fixed Issue:**

Customer Issue 113153

For large VPLEX environments, under some circumstances it is possible that discovery does not complete. As a result the supply chain can show the affected storage entities as VM datastores.

- **Fixed Issue:**

Customer Issue 113174

For vCenter Server environments, under some circumstances Turbonomic can fail to close idle sessions that it has opened with vCenter Server. This can occur after timeouts or other communication failures with vCenter.

- **Fixed Issue:**

In some cases, HTTP responses within the platform do not include the necessary security headers.

- **Fixed Issue:**

Customer Issue 113070

For Google Cloud environments, Turbonomic must discover the default CPU speed for each region.

- **Fixed Issue:**

Customer Issue 113102

Users with the Site Administrator role cannot view all the plans that have been created by other users.

- **Fixed Issue:**

When you create a reservation (to place workloads), you can define Don't Place policies for the affected scope. However, the reservation calculation ignores the Don't Place constraints.

- **Fixed Issue:**

Customer Issue 113130

In Azure environments, under some circumstances a resize can fail with the error, `Async operation failed with provisioning state: Failed: Allocation failed`.

This can happen when the resize action requires a deallocation of the initial VM instance.

- **Fixed Issue:**

Customer Issue 113099

For Nutanix environments with AHV clusters, under some circumstances resize actions fail. This can occur when Turbonomic tries to power on a VM that the Nutanix service has already powered on.

- **Fixed Issue:**

Customer Issue 113091

For Xtremlo environments, under some circumstances discovery of volumes can fail. This can occur when the volume does not have the expected entry for a volume name.

- **Fixed Issue:**

For configuration of Database policies, the user interface enables you to specify **Scaling Target VCPU Utilization** and **Scaling Target VMEM Utilization**. However, these settings are not used for Database policies, and should be removed.

- **Fixed Issue:**

Customer Issue 113079,113575

For vCenter Server environments, under rare circumstances discovery encounters a null value and fails to complete.

- **Fixed Issue:**

Customer Issue 113075

For Nutanix Prizm environments, the host capacity that Turbonomic reports can be greater than the capacity that Nutanix reports. This can occur when there are more threads on the host than cores.

- **Fixed Issue:**

If you have changed settings in the `/etc/my.cnf.d/server.cnf` file for your Turbonomic installation, when you perform an offline update the update process reverts your changes back to the defaults.

- **Fixed Issue:**

Customer Issue 113070

For Google Cloud Platform environments, under some circumstances discovery can fail. This can occur when a discovered VM is missing a required data field.

- **Fixed Issue:**

For Azure environments, under some circumstances a Service Principle can fail to return data for all of its underlying subscriptions. As a result, Turbonomic can remove the affected VMs from the supply chain.

There are occasions when the Azure environment can momentarily fail to return data for some of the underlying subscriptions. In that case, Turbonomic discovery immediately removes the affected VMs. Because this is a momentary break in data, Turbonomic should wait for a timeout period before removing the VMs from the supply chain.

- **Fixed Issue:**

Customer Issue 113054

For Kubernetes environments, under some circumstances Turbonomic does not report namespace quotas. This can occur when the quota is declared via an older syntax. This syntax is not deprecated, but it is no longer documented by Kubernetes. Turbonomic should support this older syntax.

- **Fixed Issue:**

Customer Issue 113028

For Hyper-V environments, under some circumstances Turbonomic can report excess values for storage access IOPS.

- **Fixed Issue:**

Customer Issue 113016

Under rare circumstances, for large environments with many Application Component entities, Turbonomic can generate a flood of VM Scale Up actions that it later reverts.

- **Fixed Issue:**
 When configuring schedules in an automation policy, you should not apply the same schedule to both the Policy Schedule and the Action Execution Schedule settings. If you do, then the schedule information that displays in the user interface shows that the schedule has been applied to two different policies. Also, information about the action scheduling only appears when the policy is active.
- **Fixed Issue:**
Customer Issue 113002
 For AppDynamics environments, under some circumstances during incremental discovery the AppDynamics target does not return VMem for a given VM. This can result in incorrect scaling actions.
- **Fixed Issue:**
Customer Issue 112995
 For VSan environments, if you run a plan that removes workloads, the plan can fail to suspend hosts when they are not needed. In extreme cases, this can result in more hosts than VMs in the plan scope.
- **Fixed Issue:**
Customer Issue 112988
 For VMax environments, under some circumstances the latency capacity that Turbonomic discovers can be too low. As a result, analysis can fail to generate actions that would otherwise be appropriate.
- **Fixed Issue:**
Customer Issue 112978
 In VCenter Server environments, under rare circumstances discovery can fail with the exception, `java.util.ConcurrentModificationException`.
- **Fixed Issue:**
Customer Issue 112945
 For Azure environments, if the cloud provider has a temporary outage in returning pricing data, Turbonomic charts fail to display any cost data.
- **Fixed Issue:**
Customer Issue 112922
 When inspecting a dashboard, you can display the SHOW ALL page of a given chart. If the user interface refreshes the dashboard while SHOW ALL is open, the refresh closes that page and returns to the top level of the dashboard.
- **Fixed Issue:**
Customer Issue 112875
 Under some circumstances, the user interface fails to create a new dashboard.
- **Fixed Issue:**
 If you search for VMs in the user interface, and set the filter to search for VMs with `OS = Unknown`, then the search does not return any VMs.
- **Fixed Issue:**
Customer Issue 112887
 After updating to version 7.22.9, under some circumstances analysis does not correctly assign groups to policies. This can happen for dynamic groups that are based on tags.
- **Fixed Issue:**

Customer Issue 112866

For environments that include hyperthreaded hosts, analysis can fail to scale down the VCPU for a workload that shows low utilization.

This can occur if the VM is allocated a high VCPU amount to exploit the host's hyperthreading, and that VCPU amount is greater than the host's CPU capacity.

- **Fixed Issue:**

Customer Issue 112845

When analysis is calculating actions for a group of entities that is scoped to a policy, under some circumstances the policy does not take effect for all the entities in the group, and so actions can violate the policy.

This can occur when Turbonomic fails to fully map resource utilization to the policy for one of the group members. In that case, analysis stops mapping resource utilization for the rest of the members of the group.

- **Fixed Issue:**

Customer Issue 112804,112854

Under some circumstances, when you run a plan with **Ignore Constraints** turned on, the performance is unacceptable.

- **Fixed Issue:**

Customer Issue 112821,112849,112880,113288

For vCenter Server environments, if you turn on **Enable Guest Metrics** for your target configuration, the metrics that vCenter returns to Turbonomic can be inaccurate, depending on the version of vCenter Server and VMware Tools that you have running on your VMs.

- **Fixed Issue:**

Customer Issue 112812

For Azure environments, under some circumstances an Azure target can be for a subscription that is not entered in the Azure Usage Details report. In this case, Turbonomic assumes costs for Pay-As-You-Go by default. As a result, costs can be incorrect.

- **Fixed Issue:**

Customer Issue 112794,112845

In environments that have multiple placement policies that can impact the same entity, under some circumstances analysis can generate incorrect actions. For example, analysis could generate placement actions that violate a Do Not Place policy.

- **Fixed Issue:**

For VMax environments, under some circumstances Turbonomic shows 18446744073709551615 for some monitored values. The target returns this value when there is no available data for the value, so the user interface should not display this value.

- **Fixed Issue:**

Customer Issue 112567

When you place a host in Maintenance Mode, it can take an unacceptable amount of time for Turbonomic to migrate workloads off of that host.

- **Fixed Issue:**

Customer Issue 112704,112792,112972

When running a plan that mixes VSan clusters with host clusters that do not participate in VSan, if you configure the plan to remove all the VSan hosts then the plan can fail. This can happen if the remaining VMs require the VSan storage those hosts provided.

- **Fixed Issue:**

Customer Issue 112749,113095

When you run a plan that reduces memory capacity for many VMs in a cluster, the plan simulates a reduced average utilization but does not simulate a reduced peak utilization. This can result in the plan not suspending hosts, and so it leaves the environment with many underutilized hosts.

- **Fixed Issue:**

Customer Issue 112747

When you configure a New Relic target, the Account ID field does not validate for integer-only input. This can cause confusion if you enter an invalid ID and the target fails to validate.

- **Fixed Issue:**

Customer Issue 112714

If you specify an Active Directory group for a user account, and that group name includes a period character, then you cannot subsequently delete the group from the list of user accounts.

- **Fixed Issue:**

Customer Issue 111581

For Horizon VDI environments, under some circumstances does not move a Business User to a different Desktop Pool, even though the current pool is congested. This can happen in an environment that contains many snapshot files.

- **Fixed Issue:**

For DIF developers, it is possible to create two entities of different types, and give them the same UID. For example, an Application entity and a VM entity might both use the same IP address as their UIDs. In this case, Turbonomic cannot stitch both entities into the topology, or display both entities in the user interface.

- **Fixed Issue:**

Customer Issue 112631,113016,113850

For environments that use application targets (such as SNMP) to discover application data, it is possible that the application values for VMem differ from the VMem discovered directly on the VM. In this case, the application data should take precedence. Under some circumstances, the application data does not fully replace the VM data. As a result, data display can be erroneous. For example, the 95th percentile for utilization can be lower than the average for utilization.

- **Fixed Issue:**

Customer Issue 112659,112700,113014

For Hyper-V environments that also include SQL Server targets, under some circumstances the values discovered for SQL Server can overwrite values that were discovered by the associated Hyper-V target.

- **Fixed Issue:**

Customer Issue 112593

For AWS environments, if you limit Turbonomic to only certain regions in the Region Map, it still requests spot pricing data from all regions.

- **Fixed Issue:**

Customer Issue 112590

For Azure environments, under some circumstances after a resize action succeeds, the user interface shows it as a failed action. This can occur if the communication with Azure times out when waiting for the final response to the action.

- **Fixed Issue:**

Customer Issue 112561

For AppDynamics environments, the user interface can become unresponsive.

The probe that discovers AppDynamics environments must improve its use of memory resources.

- **Fixed Issue:**

Customer Issue 112395

When you run the report *Storage Top Snapshot Consumption*, under some circumstances it can fail to list all of the snapshots in your environment.

- **Fixed Issue:**

Customer Issue 112546

Under some circumstances, when you click Show All for the Accepted Actions chart, the table only displays the first 300 actions.

- **Fixed Issue:**

For ServiceNow targets, when you configure the target to use a proxy host, under rare circumstances the target fails to validate. When validation fails, it logs the error, `Unsupported 'HTTP/1.1 407 Proxy Authentication Required'`.

- **Fixed Issue:**

Customer Issue 111995

When the Density chart is scoped to a cluster and the Show Density checkbox is selected, the displayed graph intermittently shows 0 VMs per host.

- **Fixed Issue:**

Customer Issue 112000

For Hyper-V environments where VMs use ISO storage, under some circumstances the values Turbonomic reports for storage utilization can be inconsistent between the VM and storage entities.

- **Fixed Issue:**

Customer Issue 110677

When assigning scope to a chart, for large environments the scope pick list can show only invalid group or entity types for the given chart.

This can occur because the list shows all entity types (not just valid entity types), and the list of invalid types can exceed the page length.

- **Fixed Issue:**

Customer Issue 112535

Under some circumstances, performance is unacceptable when you click Show All for the Executed Actions chart.

- **Fixed Issue:**

Customer Issue 112478

For UCS Manager environments, under some circumstances Turbonomic can lose connection with a given UCS Manager target due to timeouts.

- **Fixed Issue:**

Customer Issue 112504

For Response Time and Transactions charts, there are circumstances where the chart displays data as a percentage utilization of capacity. These charts should never show percentage values.

- **Fixed Issue:**

Customer Issue 112473

Under some circumstances, if you edit a chart after clicking in it, then Turbonomic does not save the changes to the chart.

- **Fixed Issue:**

Customer Issue 112189

For large Horizon VDI environments, under some circumstances discovery can fail to post all of the Business User entities in the supply chain.

- **Fixed Issue:**

Customer Issue 112416

When configuring a plan, if the scope includes hosts in multiple clusters, but not all of the hosts in the clusters, then under some circumstances the plan can show a current count of VMs that includes all the VMs in the affected clusters.

- **Fixed Issue:**

The performance for displaying data in Storage Summary charts and in the Storage Breakdown charts must be improved. In large environments it can take an unacceptable amount of time to get the full dataset and display it.

- **Fixed Issue:**

Customer Issue 112285

For UCS environments, under some circumstances discovery can fail when it encounters excessively large values for total throughput on a port.

- **Fixed Issue:**

Customer Issue 113204

For the Migrate to Cloud plan, when you export the Cloud Comparison chart to CSV, the resulting data does not contain more than 500 rows.

- **Fixed Issue:**

Customer Issue 112225

For the Pending Actions list, the filter includes an option to filter by Action Prerequisite. However, that filter is not supported. It should be removed from the user interface.

- **Fixed Issue:**

Customer Issue 112200

Under rare circumstances the All Actions chart can show incorrect numbers for the count of executed actions. This can occur when the Action Orchestrator component restarts while it is calculating action data.

- **Fixed Issue:**

Customer Issue 112461,112688

When configuring a plan, performance must improve for the display of the plan scope options.

- **Fixed Issue:**

Customer Issue 111889

For Azure environments, after you transfer to a new Azure Enterprise Agreement (EA), and you have set that new EA as a target, Turbonomic discovers the RI inventory that was bought under that new EA, but it no longer discovers RIs that were bought under the old EA.

- **Fixed Issue:**

Customer Issue 111668

The Target Configuration Guide should describe how to configure the Action Script target, and it should give a reference to the API guide for more information.

- **Fixed 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.

- **Fixed Issue:**

Customer Issue 110857

In a consistent scaling group, it's possible that one VM resizes for a performance reason, and other members of the group resize even though they have no performance issues. In that case, the action details show the reason to be *Efficiency*. Such a resize should be for *Compliance*.

Known Issues

- **Known Issue:**

For on-premise plans, the user interface does not display all the controls to enable or disable suspension of Host or Storage entities.

When configuring on-prem Optimize or Add VM plans, you can specify whether to enable or disable suspension of host or storage entities. However, the user interface does not include the additional controls to set the suspension to specific groups within the plan scope. You can only enable or disable suspension for the full scope of the plan.

- **Known Issue:**

Customer Issue 113954

Turbonomic does not receive pricing data for S3 database sources. As a result, actions for S3 databases can be invalid.

Due to a temporary issue with Microsoft Azure Enterprise Agreement subscriptions, Turbonomic does not receive pricing data for S3 database sources. As a result, the user interface does not show pricing data for instances of S3 databases, and actions for S3 databases can be invalid.

While this issue Enterprise Agreements persists, you should set all actions on S3 databases to MANUAL or RECOMMEND.

- **Known Issue:**

You cannot re-execute Migrate To Cloud plans that were configured on earlier releases of the platform.

Migrate to cloud plans created on Turbonomic versions 8.0.5 or earlier cannot not be re-executed. Starting with version 8.0.6 or earlier, the plan configuration requires a specification of the destination account for the migration.

- **Known Issue:**

Dashboards do not restore to original layouts after edits to charts.

After you edit a dashboard in Turbonomic, you can click to restore the original dashboard layout. However, the restore does not execute.

If this occurs, after you click to restore, refresh the browser.

- **Known Issue:**

For placement reservations, the CPU capacity in the affected cluster does not match the CPU capacity you see via external monitoring.

When you create placement reservations, the CPU capacity in the affected cluster does not match the CPU capacity you see for that cluster via external monitoring. The values that Turbonomic are off by a determined factor.

- **Known Issue:**

In the user interface, charts can incorrectly show 100% utilization of resources.

For VMM/Hyper-V 2019 environments, when Turbonomic is unable to retrieve accurate vMem utilization data for VMs for some reason, it might show 100% vMem utilization in the Capacity and Usage chart. It will, however, disable resize down actions in the background as a safeguard against incorrect actions.

- **Known Issue:**

Customer Issue 113739

Optimize Cloud plans can give inaccurate RI recommendations if you disable all actions.

If you turn on the Disable All Actions setting in the global default policy and then run an Optimize Cloud plan with VM scaling and RI purchases enabled, the plan results show inaccurate RI recommendations.

Turn off Disable All Actions to resolve this issue. Be aware that after you turn off this setting, it will take Turbonomic a week to reflect accurate results in Optimize Cloud plans.

- **Known Issue:**

For policies that exclude certain cloud tiers, when the cloud provider adds new tiers they can appear as included in the policy.

For public cloud environments, when you make a policy to include only certain tiers (VM or Storage types) of entities, if the service provider deploys new tiers then those will also be included in your policy.

This can be unexpected. For example, assume you create a policy to include only one VM type. Then if your service provider introduces new VM types, your policy will subsequently include those new types.

You should periodically check your policies to see if new tiers have been added to the INCLUDE list.

- **Known Issue:**

When you generate placement and reservations, you must be sure the associated Headroom calculations reflect changes to your environment.

When creating placement and reservations, analysis makes use of the last-generated Headroom data. This data is generated in nightly-run headroom plans. If you have changed your environment by adding targets or changing policies, then you should wait until the next run of headroom plans for the affected scope before you create placement or reservations.

- **Known Issue:**

For AWS environments, cost data is unavailable for the new EBS volume type named GP3.

For AWS environments, Turbonomic can discover the new EBS volume type named GP3. However, the cost data is unavailable at this time.

- **Known Issue:**

For environments with SNMP targets, analysis can show incorrect values for memory.

For environments that include SNMP targets, under some circumstances Turbonomic can discover incorrect memory values for Linux systems. This can occur for the systems that use the net-snmp package, version 5.7.2-43.el7. You should use versions less than 5.7.2-43.el7, or greater than or equal to 5.7.2-47.el7.

- **Known Issue:**

For Azure environments, under some circumstances analysis can repeatedly recommend and fail to execute a VM resize action.

For Azure environments, under some circumstances Turbonomic can repeatedly recommend and fail to execute a VM resize action.

Within an Azure subscription, Turbonomic can recommend resizing a VM to an instance type that is no longer supported in the given region. When this occurs, the action fails (correctly) and includes the following in the log message:

```
We do not have sufficient capacity for the requested VM size in this region.
```

This behavior is as expected. However, after encountering this error Turbonomic can repeatedly recommend the same resize action. If you encounter this situation, you can manually resize the VM to a similar, but supported instance type.

- **Known Issue:**

Customer Issue 113685

After modifying a dashboard, Reset appears to fail.

After you have modified one of the main dashboards (Application, Cloud, or On-Prem), you can click **Reset** to return the dashboard to its original configuration. However, the user interface shows that the reset process has failed.

To view the updated dashboard after performing a **Reset**, reload the browser's view of the dashboard page.

- **Known Issue:**

Changes to a policy do not immediately show up in the user interface view of the affected scope.

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 113323,113332,113493,113560

The process to change the platform's IP does not work as described in the Installation Guide.

After updating Turbonomic, the process to change the platform's IP address does not change all references to the new address.

The Installation Guide includes instructions to change the IP address of your Turbonomic installation. However, after updating to version 8.0.2 or 8.0.3, this process does not change all references to the new address. After running the script, the storage services can still show that they use the old IP address.

- **Known Issue:**

Customer Issue 113340

Hitachi Vantara targets can incorrectly show multiple compliance actions.

For Hitachi Vantara environments that enable storage replication, Turbonomic can generate repeated actions of the type, `Reconfigure Storage to Provide Extent`. In addition, replica storage entities and their providers can show incorrect utilization values.

- **Known Issue:**

When creating an automation policy, you can assign two or more schedules to the Action Execution Schedule setting. If the action is one that cannot be executed by Turbonomic, then for Service Now integrations the action appears in the Change Requests for display, only. Also, such an action is *only* affected by the first schedule in the policy.

- **Known Issue:**

For Azure environments, delays in updates to billing information can result in analysis that uses partial billing data for RIs in undiscovered subscriptions.

For the Partial RI feature, in Azure environments there can be delays in updates to billing information that Azure makes available to Turbonomic. In that case, there can be times when analysis uses partial billing data in its calculations. When showing utilization of RIs that are discovered outside of the targeted Azure subscriptions, this can result in changing costs for those RIs.

- **Known Issue:**

For Migrate to Cloud plans, under rare circumstances the plan's actions list can show duplicate entries.

For Migrate to Cloud plans, under rare circumstances the plan's actions list can show duplicate entries.

- **Known Issue:**

For Azure environments, discovery does not support the Brazil Southeast region.

For Azure environments, Turbonomic does not discover the Brazil Southeast region. Azure provides this region only to give business continuity and disaster recovery to workloads in Brazil South that require data residence.

The user interface does not display the Brazil Southeast region in any lists or charts. Also, if you do have workloads on that region, Turbonomic will not discover those workloads.

- **Known Issue:**

Customer Issue 112461**If you have configured MySQL 5.7 as an external database, you can experience poor performance when working with dynamic groups.**

If you have configured MySQL 5.7 as an external database for your Turbonomic installation, under some circumstances you can experience poor performance when working with dynamic groups. This can happen when you use extensive regular expressions as filters to generate the dynamic groups.

If you experience poor performance with dynamic groups, consider making them static groups, or consider using MariaDB as your database.

- **Known Issue:**

Customer Issue 112327**When you download the data for Pending Actions, the download might not match the data that you see in the Pending Actions chart.**

Under some circumstances, when you download the data for Pending Actions, the download does not match the data that you see in the Pending Actions chart. This can happen when the categories that the Pending Actions chart uses to group actions do not contain the correct actions. The actions are all correct, and the downloaded data groups the actions correctly.

- **Known Issue:**

When you download a CSV file from an Actions chart, the CSV file only contains the list of actions that show in the current page of data.

When you download a CSV file from an Actions chart, the CSV file only contains the list of actions that show in the current page of data. As a result, if the actions for the current scope of the chart exceed the number of entries in the page, the CSV data will be incomplete.

- **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, the platform cannot discover databases if the target authentication uses oAuth for credentials.

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

- **Known Issue:**

For existing dashboards that include the Capacity And Usage chart for databases, after an upgrade to 7.22.7 or later, the chart can appear empty.

Starting with version 7.22.7, Turbonomic tracks the DTU and Storage Amount commodities for databases. Charts that you configured for earlier versions will not include these commodities. To correct this, edit the charts to display the DTU and Storage Amount commodities. Also, when you create a new Capacity and Usage chart for databases, you must configure it to show these commodities.

- **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 user interface allows you to set orchestration for actions that the ServiceNow integration does not support.

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:**

The user interface does not currently show the billed costs for those Azure resource groups.

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 analysis 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:**

For Azure environments, analysis currently does not generate resize actions for SQL Server databases.

For Azure environments, Turbonomic currently does not generate resize actions for SQL Server databases.

- **Known Issue:**

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

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

- **Known Issue:**

Customer Issue 110123

There is a memory limit for the data you can download from the All Actions chart.

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 etcd.service can 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`

- **Known Issue:**

You must use certain templates when using PLACE to set up reservations or deployments.

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:**

For resources that do not support Reserved Capacity, charts can show them with zero reserved capacity.

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:**

Optimized Improvements for plans do not include hosts to provision.

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 high storage latency, or excessive storage provision.

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:**

The Optimal Improvements chart can show incorrect data for hosts to be suspended.

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:**

Do not edit Average or Max Utilization templates.

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:**

Do not use duplicate names for groups of the same entity type.

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 AWS payment plans that are not 1-year or 3-year terms, AWS does not return pricing data.

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 vSAN environments, under some circumstances a plan that is scoped to a datacenter can fail.

For vSAN environments, when running plans that add or replace hosts to the environment, under some circumstances the plan can show the incorrect count for hosts, and the plan can fail.

This can happen for plans that meet the following conditions:

- The plan type is Hardware Refresh, Add Workload, or Custom
- The plan scope is set to a datacenter, and it includes vSAN hosts
- The plan uses an HCI template to replace the hosts

After running, the plan shows the full count of hosts in the vSAN environment, instead of the count of hosts in the plan scope.

To avoid this situation, do not scope the plan to the datacenter.

- **Known Issue:**

Customer Issue 105693**The Headroom chart for All On-prem Hosts does not always agree with the Top Clusters chart.**

The Headroom chart for All On-prem Hosts does not always 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 vCenter Server environments, Turbonomic does not recognize DRS rules for VM restart dependencies that are based on ClusterDependencyRule.

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.