



Turbonomic 8.5.1 Release Notes

April 21, 2022

This document describes issues that are addressed in Turbonomic 8.5.1 – Release Date: April 21, 2022. 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 8.5.1, 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://support-turbonomic.force.com/TurbonomicCustomerCommunity/s/customer-support>

What's New for Version 8.5.1

Version 8.5.1

- **License Management Improvement**

The License Configuration page now displays a tabular view of your existing licenses to help you better manage these licenses. With this new view, you can sort licenses by status, expiration date, workload count, or type. You can continue to import or delete licenses as usual.

- **Real-time Analysis Improvements**

If Turbonomic real-time analysis times out, the actions that the platform generates may be incomplete. When this happens, the platform now restarts the analysis and generates a notification in Notification Center > System Health. The notification goes away as soon as the analysis completes in a succeeding market cycle. If the notification persists for an extended period of time, contact Turbonomic Support for assistance.

The default timeout is 3600 seconds (60 minutes). To use a different value, change the `rtAnalysisTimeoutSecs` property in the `CR.yaml` file.

- **Change of Support for NetApp**

Turbonomic now supports ONTAP 9.

To comply with NetApp support for ONTAP, the NetApp target no longer supports NetApp in 7-mode. For more information, please see your NetApp documentation.

Versioning Explanation

Turbonomic versioning uses V-R-M elements (Version, Release, Modification) in the version number to express the status of a given release, as follows:

Numbered Element	Example	Description
V - Version number	8.X.X	<ul style="list-style-type: none"> • Changes to platform architecture or significant changes to data models
R - Release number	X.2.X	<ul style="list-style-type: none"> • Major feature changes
M - Modification number is zero (0)	X.X.0	<ul style="list-style-type: none"> • A quarterly release • All preview features from previous bi-weekly releases are now GA • No new preview features in this release
M - Modification number is greater than zero (1 or higher)	X.X.3	<ul style="list-style-type: none"> • A bi-weekly release • Can include new preview features • Includes fixed issues

NOTE:

For API developers, the X.X.1 release can include final implementations of deprecated API features. These final implementations can make API changes that are not backward-compatible.

NOTE:

Deprecation Notice for Tbmigrate, the Classic-To-XL Migration Tool

Starting with the 8.5.0 release of Turbonomic, we will no longer support the Classic-To-XL Migration Tool. This tool provides a scripted interface that you can use to migrate from Classic installations (the 6.x version families) to XL installations (the 8.x version families).

You can still use this tool to migrate from Classic to the 8.4.x version family. We recommend that you migrate to version 8.4.0, then update to the latest version.

If you need to perform a migration after we release Turbonomic 8.5.0, you will be instructed to install and migrate to version 8.4.0, and then update that installation to the later version.

For more information, see "Appendix: Migrating Turbonomic from Classic to XL" in the *Turbonomic Installation Guide*.

Configuration Requirements

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

Turbonomic Updates and Operator Version

Turbonomic deploys as a cloud-native application on a Kubernetes cluster. This cluster can be pre-configured on a VM that you deploy, or you can deploy Turbonomic to a Kubernetes cluster in your environment. In either case, Turbonomic uses an *Operator* to manage the application deployment.

For different versions of Turbonomic, we change the version of Operator you should use, as follows:

Product Version:	Operator Version:
8.4.6	42.7
8.4.4	42.6
8.4.3	42.5
8.4.1	42.4
8.3.5	42.3
8.3.4	42.2
8.3.2	42.1
8.3.1	42.0
8.3.0	8.2

When you update Turbonomic, you should be sure to always include the matching version of Operator in the update. When you perform an online or offline update according to the instructions in the Installation Guide, the procedure should automatically include the latest Operator.

If you have installed Turbonomic on a Kubernetes cluster, you might need to manually update the Operator version.

OpenShift Installations of Turbonomic

For versions of OpenShift 4.x or later, if you manage your installation via OpenShift Operator Hub, navigate the OPC Console to **Installed Operators**. Select the version of the Turbonomic Operator that you want to use, and click **Update**. This should update the Operator, and you can now update Turbonomic to the matching version.

Other Kubernetes Installations of Turbonomic

For installations on other supported Kubernetes platforms, you can update the Operator version in either one of two ways:

- Directly edit the running deployment of Turbonomic:
 1. Enter EDIT mode for your cluster.


```
kubectl edit deployment t8c-operator -n {YourNamespace}
```
 2. Edit the Operator image.


```
Find and edit image: {required value}
```
 3. Validate that the Operator pod is ready.

Execute the command: `kubectl get pods -n {YourNamespace}`, and verify that the pod is running and ready.

- Edit the Turbonomic deployment YAML file:

1. Open the Operator deployment file for editing.

In the location where you store your manifests, open the file `operator.yaml`. This should be the file that you use to deploy the `t8c-operator` pod.

2. Edit the Operator image.

Find and edit `image: {required value}`

3. Apply the change to the operator.

```
kubectl apply -f operator.yaml
```

4. Validate that the Operator pod is ready.

Execute the command: `kubectl get pods -n {YourNamespace}`, and verify that the pod is running and ready.

After you have updated the Operator version, and you have verified that the pod is running and ready, you can edit your Custom Resource declaration to update Turbonomic to the version that matches your Operator version.

For more information, contact your support representative.

Supported MariaDB Version

For its default historical database, Turbonomic currently supports MariaDB version 10.5.13. This support includes comprehensive testing and quality control for Turbonomic usage of the historical database.

IMPORTANT:

Do not update MariaDB to later versions (10.5.14 or 10.5.15). Turbonomic does not support operation with versions later than 10.5.13.

If you are running Turbonomic installed as an OVA, and using the database that is included in that OVA installation, then you must use version 10.5.13. For versions of Turbonomic that you installed as an OVA *before* version 8.4.6, it is likely that you need to update your MariaDB.

For information about updating your MariaDB instance, see "Verifying your MariaDB Version" in the latest version of the *Turbonomic Installation Guide*.

Required DB Capacities

For Turbonomic versions 8.0.6 or later, your historical database must provide certain storage size capacities. For MariaDB or MySQL installations, you must ensure your database provides the necessary messaging and logging capacity.

If you installed Turbonomic as an OVA, and use the included MariaDB for the historic database, the easiest way to set the correct capacities is to update your Turbonomic to version 8.0.7 or later. For more information, see "Increasing your Database Capacities" in the latest version of the *Turbonomic Installation Guide*.

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:

```
{ {ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION} }
```

In particular, the SQL modes should *not* include `ONLY_FULL_GROUP_BY`, `NO_ZERO_IN_DATE`, or `NO_ZERO_DATE`.

External DBs and Turbonomic Updates

If you have deployed Turbonomic with an external database server, for some updates you might need to manually create a new database and user for that deployment. This is important if your external database server is multi-tenant, or if your deployment does not grant administrative privileges to Turbonomic.

NOTE:

If your external database server is multi-tenant, or if your database server does not grant administrative privileges to Turbonomic, then you must continue with this configuration requirement.

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

If you deployed your database server in a way that grants Turbonomic privileges to create new databases and new users, then the update will automatically create the required database. This configuration requirement does not apply to you and you do not need to take any action.

For some Turbonomic updates, the updated version includes new databases on the historical database server. If you are updating to one of these versions, then you must *first* create the new database, and a user account with privileges to access that database. After you have done this, you can update to the latest version of Turbonomic.

For more information, see "External DBs and Turbonomic Updates" in the latest version of the *Turbonomic Installation Guide*.

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 TLS 1.2 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 TLS 1.0. 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.

Fixed Issues

- **Fixed Issue:**

Customer Issue 119691

Under some circumstances a webhook workflow can execute twice.

If you configure a webhook workflow to replace a given action, under some circumstances the webhook can execute twice for a single action.

- **Fixed Issue:**

Customer Issue 119400

For Instana targets, under rare circumstances target validation can fail.

When you configure an Instana target, under rare circumstances target validation can fail to complete.

- **Fixed Issue:**

Customer Issue 119594

Downloaded CSV data from the Top Volumes chart is confusing.

When you download CSV data from the Top Volumes chart, it shows storage in MB when it should be in GB. Also, cost values can be incorrect.

- **Fixed Issue:**

Under rare circumstances, action analysis can fail.

Under rare circumstances action analysis can fail with the following log message:

```
Need durable recommendation OID for action
```

- **Fixed Issue:**

Customer Issue 119537

For VDI environments, analysis does not generate resize down actions for logged-off Business User entities.

For VDI environments, analysis does not generate resize down actions for logged-off Business User entities.

- **Fixed Issue:**

Customer Issue 119427

For Azure environments, under some circumstances discovery of the EA target can fail.

For Azure environments, if the `UsageDetailsReport.csv` file includes a malformed tag, then discovery of the EA target can fail.

- **Fixed Issue:**

Customer Issue 119458

For AWS environments with billing accounts but no master account, you cannot add a discount for the billing family.

For AWS environments with billing accounts but no master account, the user interface to configure price adjustments does not display the billing account. As a result, you cannot configure a discount.

- **Fixed Issue:**

Customer Issue 119424

For Azure, under some circumstances a Delete Volume action can fail.

For Azure environments, the Delete Volume action fails to execute if the volume name ends with `.vhd`

.

- **Fixed Issue:**

Customer Issue 119224

When creating reservations that include additional storage, analysis can fail to place the new storage.

If you create a reservation that is scoped to a specific cluster, and the reservation adds storage capacity, the reservation analysis can fail to place the storage.

- **Fixed Issue:**

Customer Issue 119274

When you create groups, the filter to exclude resource groups can fail to exclude the indicated entities.

In Azure environments, you can create dynamic groups that use a filter to exclude members of a given resource group (filter is **Resource Group != {RG_NAME}**). However, if you add filters to exclude more than one resource group, the filters do not exclude any entities.

- **Fixed Issue:**

Customer Issue 119055

For Horizon VDI environments, under some circumstances discovery can fail.

When discovering a Horizon VDI environment, analysis executes queries at different times to collect the full range of data. If a VM found in an initial query has been removed, and the subsequent query does not find that VM, discovery can fail.

- **Fixed Issue:**

Customer Issue 119059,119176,119621

Under rare circumstances, the user interface can stop showing discovered metrics.

Under rare circumstances, the process that saves newly discovered statistics to the database can lock up, and Turbonomic can stop saving incrementally discovered metric data. This can result in gaps in timeline data, or plans can time out.

- **Fixed Issue:**

Customer Issue 118954

For some Azure subscriptions, discovery can take an excessive amount of time.

For some Azure subscriptions, discovery can take many hours to complete. This can occur as discovery requests metric tables from the Azure environment, especially if the subscription has storage a large number of these tables.

- **Fixed Issue:**

Customer Issue 119014,119347,119404

For Azure, if execution of an action fails, that can cause subsequent action execution to fail.

For Azure environments, if a execution of an action fails, under some circumstances analysis can fail to free the action lock. As a result, subsequent actions will also fail to execute.

Security Issues

At Turbonomic we do our best to ensure Turbonomic is a product you can fully trust to manage the performance of your enterprise applications. Delivering a secure product is an important part of gaining that trust. We know our products must meet the highest security standards so you can deploy them with confidence. This section lists the vulnerabilities we have identified and corrected.

If you identify a potential security vulnerability, please report it to your Technical Support representative as soon as possible.

- **Fixed Issue:**

Security Update: com.sun.xml.ws:jaxws-rt

This release updates the version of com.sun.xml.ws:jaxws-rt to address a known security issue.

- **Fixed Issue:**

Security Update: org.postgresql:postgresql

This release updates the version of org.postgresql:postgresql to address the following issue:

- CVE-2022-262520

- **Fixed Issue:**

Security Update: gson

This release updates the version of gson to address a known security issue.

- **Fixed Issue:**

Security Update: xmlsec

This release updates the version of xmlsec to address the following issues:

- CVE-2019-12400
- CVE-2021-40690

- **Fixed Issue:**

Security Update: Eclipse Equinox

This release updates the version of Eclipse Equinox to address the following issue:

- CVE-2021-41033

- **Fixed Issue:**

Security Update: commons-collections

This release updates the version of commons-collections to address the following issues:

- CVE-2015-6420
- CVE-2017-15708
- **Fixed Issue:**

Security Update: jackson-databind

This release updates the version of jackson-databind to address the following issues:

- CVE-2020-36518
- **Fixed Issue:**

Security Update: tomcat-api

This release updates the version of tomcat-api to address the following issue:

- CVE-2022-23181
- **Fixed Issue:**

Security Update: Spring

This release updates the version of Spring to address the following issues:

- CVE-2022-22965
- CVE-2022-22963
- **Fixed Issue:**

Customer Issue 119435**Security Update: OVA**

This release updates the OVA to address the following issue:

- CVE-2016-2183

Known Issues

- **Known Issue:**

Filters for search and groups based on target type do not support Azure Subscription or GCP Project.

For VMs, when searching or creating a group based on target type, the user interface does not include options to filter by **Azure Subscription** or **GCP Project**. At this time you cannot create a filter on these criteria.

- **Known Issue:**

Customer Issue 119471,119789

For Headroom plans, you can specify the desired state for overprovisioned commodities.

Turbonomic runs nightly plans to calculate headroom in clusters - The number of VMs you can add to a cluster without exceeding limits for Memory, CPU, or storage. For example, if you want a desired state of 80% consumption, the plan will not add VMs that will exceed 80% utilization of resources in the cluster.

Note that reservations use these headroom calculations to determine whether Turbonomic can place the workloads in a reservation request.

By default, plans do not keep the utilization of overprovisioned resources within the desired state (80% in the above example). The plan calculates 100% utilization of overprovisioned resources. However, when placing reservations in the cluster, analysis could recommend provisioning new hosts as a way to keep the overprovisioned resources within the desired state.

We have introduced a setting you can make to enforce the desired state on overprovisioned resources. This can result in a lower calculated headroom. But when a reservation places and powers on a new VM, it is guaranteed to fit on the current infrastructure. Analysis will not see a need to provision a new host to support the VM's overprovisioned resources.

To turn on this feature, edit the topology-processor settings in the cr.yaml file:

1. Open an SSH terminal session on your Turbonomic instance.

Log in with the System Administrator that you set up when you installed Turbonomic.

2. Open the cr.yaml file for editing. For example:

```
vi /opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml
```

3. In the `spec/properties` section, find the entry for the `topology-processor` component.

4. Add the following properties to the component spec:

```
considerDesiredStateForProvisioningInClusterHeadroomPlan: true
considerUtilizationConstraintInClusterHeadroomPlan: true
```

5. Save and apply your changes to the platform.

After you save your changes, use `kubectl` to apply the changes:

```
kubectl apply -f /opt/turbonomic/kubernetes/operator/deploy/crds/
charts_v1alpha1_xl_cr.yaml
```

For assistance, contact your Support representative.

- **Known Issue:**

For Kubernetes, in some environments analysis cannot execute Scale Node actions.

For Kubernetes OCP 4.x and AKS environments, Turbonomic can generate and execute Scale Node actions. However, if the environment includes other K8s distributions that don't support execution of Scale Node actions (EKS, AKS, and OCP), then Turbonomic can disable execution of *all* Scale Node actions in the environment.

- **Known Issue:**

For AWS, discovery does not recognize db.m6i or db.r6i workload tiers.

For AWS environments, discovery does not recognize the db.m6i or db.r6i workload tiers. If you have AWS managed databases from those tiers in your environment, discovery will skip them, and you will see an error similar to:

```
Cannot find attributes for :db.m6i.8xlarge...
```

- **Known Issue:**

For Kubernetes environments with nodes running Linux with cgroup v2 enabled, you must use Kubernetes version 1.23.2 or later.

For Kubernetes environments with nodes running Linux with cgroup v2 enabled, an issue in earlier versions of Kubernetes prevents Turbonomic from collecting CPU utilization data for the affected nodes. To collect CPU utilization from the cgroup v2 nodes, you must run Kubernetes version 1.23.2 or later.

- **Known Issue:**

Customer Issue 118620,119176

For very large environments, under rare circumstances queries can lock the database and cause gaps in discovered data.

For very large environments, it is possible for queries to lock the database, and cause gaps in discovered data. When this occurs you can see errors similar to the following:

- [RollupProcessor] : Error during rollup activity for table vm_stats_latest:
- [ComponentBasedTargetDumpingSettings] : Retaining 0 discovery dumps for target

If this occurs in your environment, contact your support representative for a workaround.

- **Known Issue:**

Billed costs shown in Turbonomic and Google Cloud Platform (GCP) do not match due to time differences.

Billed costs shown in Turbonomic charts and GCP billing reports do not match because Turbonomic uses UTC, while GCP uses local time. However, costs shown in both places are correct and reliable.

- **Known Issue:**

For IBM FlashSystem, a known FlashSystem issue can cause some models to show intermittent target errors.

For some models of IBM FlashSystem platforms, Turbonomic can show intermittent `Invalid Credentials` target errors. This is caused by a known FlashSystem REST service issue with platforms that have less than 64GB of memory.

If you experience this issue, use the following command to restart the FlashSystem REST service:

```
satask restartservice -service cfrest
```

Please contact your IBM FlashSystem Support representative for more information.

- **Known Issue:**

Customer Issue 117905

For workloads with very large disks, Storage vMotion can time out.

Storage vMotion actions for VMs with very large disks can time out. If you experience this, contact your support representative for help changing the timeout threshold.

- **Known Issue:**

Customer Issue 117307

When updating your installation, under rare circumstances the Topology Processor component can fail to restart.

When updating Turbonomic, under rare circumstances the `topology-processor` pod can fail to restart. The log posts an error with the statement:

```
AccessDeniedException: /home/turbonomic/data/kv
```

If you experience this problem, please contact your support representative for a workaround.

- **Known Issue:**

For Azure, targets in the Australia Central region can show inconsistent pricing for VM license costs.

For Azure environments running in the Australia Central region, the pricing for license costs that is reported to Turbonomic can be incorrect. As a result, the user interface can display incorrect license costs for Linux or Windows operating system licenses.

- **Known Issue:**

For Azure and AWS, analysis does not consider the full cost for some workloads.

For Azure environments, analysis considers the base OS cost, but does not consider additional costs for support or other add-on features that are bundled with the OS. The affected OS types are Ubuntu PRO, SUSE 24/7, and RHEL with HA.

For AWS environments, analysis does not consider AWS Marketplace costs.

- **Known Issue:**

For certain installation configurations, an online update can take longer than expected.

If you have installed Turbonomic behind a firewall and opened ports for the Docker Hub services that deliver the Turbonomic, it is possible that the script `onlineUpgrade.sh` can take 30 to 45 minutes to complete. If you have trouble running the update, please contact your support representative.

- **Known Issue:**

After updating the platform, Embedded Reports can fail to display.

Under some circumstances, after you update Turbonomic to a new version, the Embedded Reports page can fail to display. The update appears to have completed successfully, but the Embedded Reports components appear to be running and ready. However, when you click on the Embedded Reports button, you get the following error:

```
failed to log in as user, specified in auth proxy header.
```

This can occur when the host VM takes a long time to completely start up all the platform components. If you encounter this problem, make sure all the components are running, and then execute the following command to restart the `grafana` pod:

```
kubectl delete pod -l app=grafana
```

For assistance, contact your support representative.

- **Known Issue:**

For Citrix XenCenter environments with VMs that share storage, analysis can report excess capacity for thin-provisioned storage.

For Citrix Hypervisor targets, analysis can calculate thin-provisioned capacity for storage. However, if multiple VMs share that storage, the calculated storage capacity can be doubled for each additional VM.

- **Known Issue:**

Customer Issue 116930

For Fabric environments, when you scope the view to Datacenter the supply chain can fail to include all the related Host entities.

For environments that include Fabric targets, under some circumstances a view that you scope to the Datacenter might not include all the related Host entities. This can happen for hosts that include the hyphen character ("-") in the Host name. For example, for Cisco UCS targets, the supply chain will not show UCS hosts (that use the hyphen character in their names) when you scope to the Datacenter entity.

- **Known Issue:**

Customer Issue 115635

For Azure, under some circumstances a successful scaling action appears in the logs as failed.

For Azure environments that include Availability Sets, under some circumstances a successful scaling action appears in the Turbonomic logs as failed.

- **Known Issue:**

Customer Issue 114942,118692

With Executed Actions charts, some data is missing for actions on entities that have been removed from the environment.

When you view Executed Actions charts or export data from them, some data is missing for actions on entities that have been removed from the environment. For example, assume an action was executed on a storage volume, and that volume has later been removed from the environment. In that case, the exported data for that action will not include values that describe the removed volume.

- **Known Issue:**

For Azure, if you connect the target through a proxy, then the target does not discover unattached storage volumes.

For Azure environments, if you connect to the Azure target through a proxy, then Turbonomic does not discover unattached volumes.

This issue occurs because the AzureVolumes probe must update the version of the azure-storage-sdk that it uses.

- **Known Issue:**

The Onboarding wizards can sometimes fail to close.

When you first install Turbonomic, the user interface displays onboarding wizards to walk you through setting up your license, and configuring your first target. Under some circumstances, the button to end the wizard's workflow does not close the wizard. That can result in blocking you from continuing your Turbonomic session.

If the onboarding wizard does not close when you click **End Setup**, refresh the browser. That should close the wizard and leave you on the last user interface page that you visited.

- **Known Issue:**

For Kubernetes environments, when you enable Feedback and Diagnostics for your installation, the collected data can include Kubernetes cluster names.

To help us improve the product, you can enable Turbonomic to collect anonymized and non-confidential data as you go about using the product. However, because of the way Kubernetes discovery works for Turbonomic, the collected data includes the names of any Kubernetes clusters that you have set up as targets. We do not make use of those cluster names in any way.

If you do not want Turbonomic to collect these cluster names, then you can navigate to **Settings / Maintenance Options / Feedback and Diagnostics** and turn off the option to share anonymized usage data.

- **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 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 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:**

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 Database Servers if the target authentication uses OAuth for credentials.

For AppDynamics environments, Turbonomic cannot discover Database Servers if the target authentication uses OAuth for credentials.

- **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 some 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:**

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:**

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.