



Turbonomic 8.2.1 Release Notes

June 2, 2021

This document describes issues that are addressed in Turbonomic 8.2.1 – Release Date: June 2, 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 8.2.1, 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 8.2.1

Version 8.2.1

- **AWS Savings Plans Utilization**

Use the **Savings Plans Utilization** chart to evaluate how well you have utilized your [AWS Savings Plans](#) inventory. Turbonomic calculates utilization based on the billing data that it discovers via AWS Cost Explorer.

NOTE:

For cloud environments with both Reserved Instances and Savings Plans, utilization data is included in the **RI and Savings Plans Utilization** chart.

For details, see *Savings Plans Utilization* in the User Guide.

Configuration Requirements

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

Supported MariaDB Version

For its default historical database, Turbonomic currently supports MariaDB version 10.5.6. This support includes comprehensive testing and quality control for Turbonomic usage of the historical database.

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 10.5.6*. For versions of Turbonomic that you installed as an OVA *before* version 8.0.6, 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 greater. 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.

This table lists the Turbonomic versions that require new databases. If you are updating from an earlier version to one of these, you must create the indicated new databases. For example, if you are updating from version 8.0.0 to 8.1.1, then you must create the `repository` database and the `market` database.

Turbonomic Version:	New Databases:	Notes:
8.0.1	<code>repository</code>	If you are updating from a version earlier than 8.0.1, you must create a database named <code>repository</code> , and a user account named <code>repository</code> .
8.1.1	<code>market</code>	If you are updating from a version earlier than 8.1.1, you must create a new database named <code>market</code> , and a user account named <code>market</code> .

NOTE:

If you have already updated to one of these versions of Turbonomic, and you did not perform the steps to update your external DB, please contact your support representative.

To create the databases and users, you will:

- Manually create each required database

This includes creating the database in your DB instance, creating a user to access the database, and granting privileges to the user.
- Manually add each required database to your `cr.yaml` file

The `cr.yaml` file declares entries for each component database. 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 each new database.

To create a new database:

1. Connect to your external 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, where `<New_Database>` matches the database name in the table above:

```
create database <New_Database>;
```

For example, to create a new `market` database, execute:

```
create database market;
```

3. Create the account that Turbonomic will use to access the database where `<New_Database>` matches the database name in the table above:

```
create user '<New_Database>'@'%' identified by 'vmturbo';
```

For example, to create a user for the `market` database, execute:

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

NOTE:

The value `vmturbo` is the default password that Turbonomic 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 new user account, where `<New_Database>` matches the database name in the table above:

```
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 <New_Database>.* TO '<New_Database>'@'%' ;
```

For example, to set account privileges for the `market` user, execute:

```
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 market.* TO 'market'@'%' ;
```

5. Flush privileges to make the privileges take effect:

```
flush privileges;
```

Now that the new database is created in your external DB service, you must declare access to it 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 matching pod to access the new database.

Add the entry to the `properties:` section of the `cr.yaml` file, where `yourPassword` is the password that you assigned to that user account, `yourDBServer` is the qualified name of your external DB or your multi-tenant DB partition, and `<New_Database>` is the name of your new database. Declare the following in the entry you add:

```

<New_Database>:
  <New_Database>DbUsername: <New_Database>@yourDBServer
  <New_Database>DbPassword: yourPassword
    
```

For example, if you added the `market` database, the resulting `cr.yaml` file should be similar to the following. (Note that the `global` properties assume that you want SSL connections, and the database is MySQL):

```

properties:
  global:
    enableSecureDBConnection: true
    sqlDialect: MYSQL
    dbRootUsername: dbAdminUser@yourDBServer
    dbRootPassword: dbAdminUserPassword
  action-orchestrator:
    actionDbUsername: action@yourDBServer
    actionDbPassword: yourPassword
  auth:
    authDbUsername: auth@yourDBServer
    authDbPassword: yourPassword
  clustermgr:
    clustermgrDbUsername: clustermgr@yourDBServer
    clustermgrDbPassword: yourPassword
  cost:
    costDbUsername: cost@yourDBServer
    costDbPassword: yourPassword
  group:
    groupComponentDbUsername: group_component@yourDBServer
    groupComponentDbPassword: yourPassword
  history:
    historyDbUsername: history@yourDBServer
    historyDbPassword: yourPassword
  plan-orchestrator:
    planDbUsername: plan@yourDBServer
    planDbPassword: yourPassword
  topology-processor:
    topologyProcessorDbUsername: topology_processor@yourDBServer
    topologyProcessorDbPassword: yourPassword
  repository:
    repositoryDbUsername: repository@yourDBServer
    repositoryDbPassword: yourPassword
  market:
    marketDbUsername: market@yourDBServer
    marketDbPassword: yourPassword
    
```

After you have done this, you can update to the latest version of Turbonomic. (Note that upgrading applies changes to the version information in this 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 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 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 115503

Exported data for an on-prem Delete Volume action now limits the list of related entities for that action.

The Data Export service includes related entities with the data for delete actions. However, in previous versions it included too many related entities for a Delete Volume in an on-prem environment. As a result, it was difficult to filter the data by related entity.

Starting with this version, the Data Exporter limits the list of related entities to the following types:

- Logical Pool
- Disk Array
- Storage Controller
- Host
- Datacenter
- Group (normal groups containing entities above)
- Cluster
- Storage Cluster

- **Improvement:**

Customer Issue 115555

Policy settings for Latency and IOPS capacity now override the discovered capacities.

For Logical Pool/Disk Array targets (for example EMC VMAX), Turbonomic discovers Latency and IOPS capacities that it uses as defaults. Starting with this release, if you override those defaults via policy settings, analysis will use the override settings that you make in those policies.

- **Improvement:**

Customer Issue 112776

For Migrate to Cloud Plans, the downloaded data now includes costs as numbers, not as text.

When you download data for a Migrate to Cloud plan, the Before and After Cost values are now formatted as numeric values, and not as strings.

Fixed Issues

- **Fixed Issue:**

When you download data from the Buy RIs action center, the download menu does not display the correct text.

When you download data from the Buy RIs action center, the download menu displays a command similar to `Buy_displayname`. The menu command should read `Buy Reserved Instance`.

- **Fixed Issue:**

Customer Issue 115830

The download menu for the Recommended RI Purchases chart does not display any commands.

When you try to download the data from the **Recommended RI Purchases** chart, the download menu is blank. It does not display any commands, so you cannot execute the download.

- **Fixed Issue:**

Customer Issue 115816

When using the SEARCH view to find groups, the view can fail to show groups that exist in the environment.

When using the SEARCH view to find groups, the view can fail to show groups that exist in the environment. However, if you search via filters and use a Regex expression to find the group, then Search can find it.

- **Fixed Issue:**

Customer Issue 115761

Under some circumstances discovery can fail to connect Kubernetes entities with vCenter Server entities.

For Kubernetes environments that include vCenter Server, under some circumstances Turbonomic can fail to stitch vCenter entities in the Kubernetes supply chain. This can occur when the pod specification includes duplicate Affinity rules.

- **Fixed Issue:**

Customer Issue 115485

For Embedded Reports, the Wasted Files report does not include an important warning.

For Embedded Reports, the Wasted Files report should include the following warning in its header:

“WARNING: The information in this report is correct as of the time the data was collected. The data can be up to one day older than the date of report generation. Before you delete files based on this report, verify that the files are indeed unused.”

- **Fixed Issue:**

Customer Issue 115687

When you download data from the new Actions Center View for Volume Scaling actions, the data is not complete.

If you enable the **New View** for the Actions Center, the download for Volume Scaling actions data does not include the `Linked VMs` column. To match the older view, this data should include that column.

- **Fixed Issue:**

Customer Issue 115588,115677

Under some circumstances an Extractor report editor user can lose privileges that the user would have with Embedded Reports.

Under some circumstances an Extractor report editor user can lose privileges that the user would have with Embedded Reports. This can happen if the license processing in Turbonomic encounters an error.

- **Fixed Issue:**

Customer Issue 115547

For some Kubernetes environments the user interface can fail to show container CPU and Memory data.

For some Kubernetes environments the user interface can fail to show container CPU and Memory data. This can occur because Kubernetes has changed some of the metric names that Turbonomic analysis must use.

- **Fixed Issue:**

Customer Issue 115652

For AWS, under some circumstances analysis can fail when parsing cost data.

For AWS environments, under some circumstances analysis can fail when parsing cost data. This can occur when the parser encounters unexpected entries in the data.

- **Fixed Issue:**

Customer Issue 115598

Multiple resize actions on a single VM execute in separate phases, which can keep a VM powered down longer than necessary.

There are circumstances where Turbonomic recommends more than one resize action for the same VM. For example, it could recommend resizing VCPU and also VMEM for the same VM. In that case, Turbonomic executes the actions, one after the other. This would take a minimum of two 10-minute cycles.

- **Fixed Issue:**

Customer Issue 115455

For Migrate to Cloud plans, under some circumstances the plan incorrectly includes migration actions for storage volumes.

For an on-prem storage, if Turbonomic does not discover a volume for that storage, then the Migrate to Cloud plan can include that storage in the migration. For example, this can occur for an ISO or NFS share that Turbonomic models as a storage.

- **Fixed Issue:**

Customer Issue 115263

For datastores with duplicate names, under some circumstances the product only discovers one of the datastore instances.

In one environment there can be multiple datastores with the same server IP, sane remote path, and same name. For example, you can use the same naming for two datastores in different datacenters. In that case, discovery treats these datastores as a single instance.

- **Fixed Issue:**

Customer Issue 115228

Under some circumstances a download from the Action Center can fail.

When you download data from the Action Center, the download can fail if the process encounters undefined values for the storage amount.

- **Fixed Issue:**

Customer Issue 114354

If you have enabled Feedback and Diagnostics, but your system cannot access DataCloud, then the Prometheus pod can have containers that are not in a ready state.

When the Prometheus pod fails to connect with DataCloud, two of its containers can remain in a state waiting for the DataCloud data. This does not affect Turbonomic operation.

- **Fixed Issue:**

Customer Issue 114004

For WMI environments, under some circumstances the target can fail and requires a manual restart of the probe component.

After a WMI probe fails, it should restart automatically.

Note that because of authentication constraints, it can take up to an hour to automatically restart a failed component. If you must restart it sooner than that, you can manually restart it.

Known Issues

- **Known Issue:**

Updates from version 7.22.6 to version 8.2.0 can fail to complete.

If you are updating Turbonomic from version 7.22.6 to version 8.2.0, *and* you have enabled Embedded Reporting, then the update can fail to complete. If this occurs, recover the snapshot of your original installation and try to update to an earlier version (for example, 8.1.6). Then try to update that version to 8.2.0.

If the problem persists, please contact your support representative.

- **Known Issue:**

If you update from versions 8.1.4, 8.1.5, or 8.1.6, then Embedded Reports can fail to connect.

NOTE:

If you *initially installed* Turbonomic as version 8.1.4, 8.1.5, or 8.1.6, this note applies to you.

If you have enabled Embedded Reports for your installation, then update to version 8.2.0 or later, under some circumstances the Embedded Reports page can fail to load when you click **Reports**. If this occurs, it is because the `grafana` pod failed to start up and enter the ready state.

If you experience this, you will not lose any data, and the Embedded Reports database still collects the Turbonomic data as usual. After you restart the `grafana` pod, your reports will show continuity of data.

To regain connection to the Reports page, restart the `grafana` pod:

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

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

- Username: `turbo`
- Username: `[your_private_password]`

2. Restart the `grafana` pod.

```
kubectl scale --replicas=0 deployment grafana
```

3. Verify that the `grafana` pod is running.

Give the platform enough time to restart the pod, then execute the command:

```
kubectl get pods -n turbonomic
```

You should see output similar to the following:

NAME	READY	STATUS	RESTARTS
...			
grafana-5f41dd61c4-4d61q	1/1	Running	0
...			

- **Known Issue:**

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 when you scope to the Datacenter entity.

- **Known Issue:**

After you change which user has Report Editor privileges, it can take 30 minutes for the change to take full effect.

For Embedded Reports, you can specify which Turbonomic user has Report Editor privileges. If you grant these privileges to a different user, it can take up to 30 minutes before the Reports page will show the Report Editor username.

- **Known Issue:**

Customer Issue 115352,115394,115453

Updates to the product should always use the latest Operator.

When you update Turbonomic you should be sure to always include the latest 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.

A symptom you might see if you do not have the latest Operator is that the Embedded Reports component (the `grafana` pod) fails to start up.

If you need to install the latest Operator manually, you can execute:

```
sudo docker pull turbonomic/t8c-operator:8.2
```

- **Known Issue:**

Customer Issue 115153,115439

- **Known Issue:**

Customer Issue 114942

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:**

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 work flow does not close the wizard. That can result in blocking you from continuing on with your Turbonomic session.

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

- **Known Issue:**

For Hyper-V, Hardware Replace plans that use HCI Host templates can give inconsistent results.

For Hyper-V environments, if you run a Hardware Replace plan that replaces hosts with HCI Host templates, the results can be inconsistent. Under some circumstances the plan can fail to place all the VMs in the plan scope.

- **Known Issue:**

For plans that were run in earlier versions, the plans cannot load Azure RI Buy actions when you view them in later versions.

For plans that were run in versions 8.0.7 or earlier, the plans cannot load Azure RI Buy actions when you view them in later versions.

Because of anomalies that can appear in the RI pricing data that Azure sends to Turbonomic, it was necessary to change the pricing data that analysis uses. This has no impact on real-time analysis. However, if you load a plan that was run in an earlier version, and that plan includes Buy RI actions, then those actions will not appear in the plan.

- **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:**

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:**

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:**

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 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:**

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:**

Customer Issue 107699,108212,109349,110474,110896,112196,112639,112749,112790,113104,114386,115748

Hardware Refresh plans to replace hosts can give unexpected results.

When you run a plan to replace hosts in a cluster, the results can incorrectly show that you need more hosts than you would expect. This can occur for two reasons:

- If the host template for the replacements does not use CPU specifications from the catalog, then the calculations for host capacity can be incorrect.
- When placing VMs on the replacement hosts, analysis assumes all the VM peaks can occur at the same time. This results in excessive utilization of peaks capacity in the plan.

To configure a plan in a way that avoids these problems, you can:

- Make sure the plan scope is for a single cluster.

This is the typical use case for a Replace Hosts plan. Record the cluster name, because the name will come in handy in later steps.

- Always use the **Select from Catalog** option when you create your Host template.

If the catalog does not include the CPU specifications that you want, choose an entry that is as close as possible.

To adjust the number of cores, you can then specify a different number of sockets. For example, assume you choose a CPU spec that includes 4 cores of a given core frequency, but you actually want 32 cores. You can choose that spec, and then set **Sockets** to 8 to achieve that number.

- Turn off **Scale** for the VMs.

After you choose the host template to use, click **NEXT: VIRTUAL MACHINE ACTIONS**. Then turn off the **Scale** option.

When replacing hosts, it's important to *not* scale the VMs, so you can see how the hosts can support your given workload.

- Replace all your VMs with a Cluster Average template.

Click **SKIP TO CONFIGURATION** to display the plan configuration, and open **Replace/Virtual Machine**. Display **Clusters** and click **Select all**. Then click **NEXT** to choose the VM template.

To choose the VM template you will replace with, type your cluster name in the Search box. The Templates list will show an AVG template for that cluster. For example, if your cluster name is MyCluster1, then the template name would be `myDomain.com: :AVG:MyCluster1`. This template captures the average VM utilization over the last ten days.

Choose this template and click **SUBMIT**.

- You can now run the plan.

- **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.