



Turbonomic 7.22.1 Release Notes

May 26, 2020

This document describes issues that are addressed in Turbonomic 7.22.1 – Release Date: May 26, 2020. Please see the Turbonomic 7 documentation for earlier versions of the Release Notes:

<https://greencircle.vmturbo.com/community/products/pages/documentation>

For any questions, please contact Turbonomic Technical Support at support@turbonomic.com, or open a ticket at:

<https://greencircle.vmturbo.com/support>

What's New for Version 7.22.1

This release of Turbonomic includes the following new features:

- 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.

- 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.

- **EARLY ACCESS FEATURES:**

This release introduces previews for the following features. We have performed limited testing of these features, but we encourage you to try them in controlled settings, and provide us with any feedback you might have.

- Improved Container Management

For Kubernetes environments, this release introduces new entity types to represent your environment in the supply chain. Turbonomic now discovers:

- Namespace
- Workload Controller
- Container Spec

Introducing these entity types in the model enables analysis to better reflect the utilization and requirements of your environment, which results in better actions to maximize performance and savings.

This is an Early Access release of the new Kubernetes support. To discover these entity types, Turbonomic introduces a new version of the Kubeturbo service. For information about the latest Kubeturbo releases, go to <https://github.com/turbonomic/kubeturbo/wiki/Server-Versions-and-Kubeturbo-Tag-Mappings>

- Integration with VMware Horizon VDI

For server-based deployments of Turbonomic, only:

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
- Desktop Pool – A collection of desktops that support users
- View Pod – A logical grouping of desktop pools, that provides a capacity for active sessions

- Improved vSAN Support for Hyperconverged Environments

For server-based deployments of Turbonomic, only:

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.

- Nutanix Support

For server-based deployments of Turbonomic, only:

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.

Configuration Requirements

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

Transport Layer Security Requirements

By default Turbonomic requires Transport Layer Security (TLS) version 1.2 to establish secure communications with targets. Most targets should have TLSv1.2 enabled. However, some targets might not have TLS enabled, or they might

have enabled an earlier version. In that case, you will see handshake errors when Turbonomic tries to connect with the target service. When you go to the Target Configuration view, you will see a Validation Failed status for such targets.

In particular, we have found that NetApp filers often have TLS disabled by default, and that the latest version they support is TLSv1. If your NetApp target fails to validate, this is could be the cause.

If target validation fails because of TLS support, you might see validation errors with the following strings:

- No appropriate protocol

To correct this error, ensure that you have enabled the latest version of TLS that your target technology supports. If this does not resolve the issue, please contact Technical Support.

- Certificates does not conform to algorithm constraints

To correct this error, refer to the documentation for your target technology (for example, refer to NetApp documentation) for instructions to generate a certification key with a length of 1024 or greater on your target server. If this does not resolve the issue, please contact Turbonomic Technical Support.

Improvements

- **Improvement:** For deployments that use Single Sign On (SSO), with this release Turbonomic can 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.
- **Customer Issue 110646,110931**

Improvement: Consistent Scaling policies for Discovered Groups

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, the consistent scaling was set to these groups silently, and there was no representation of that setting in the user interface.

With this release, Turbonomic discovers the groups and creates read-only policies for them to implement the 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**.

The entities that Turbonomic puts into consistent resize groups are:

- Azure Scale Sets
- AWS Autoscaling Groups
- Kubernetes Groups: Deployments, ReplicationControllers, ReplicaSets, DaemonSets, and StatefulSets

Also note that for Azure Availability Sets, consistent resizing is not turned on. However, Turbonomic does create a policy that identifies compute tiers to exclude from that set.

Fixed Issues

- **Customer Issue 110846**

Fixed: Performance to display data in the Storage Summary chart must be improved.

- **Customer Issue 110942**

Fixed: For Google Cloud Platform environments that include multiple targets, under some circumstances Turbonomic does not discover an accurate supply chain of entities.

- **Customer Issue 110928**

Fixed: For public cloud environments with RIs, under some circumstances if there are multiple instances of an RI, analysis calculates incorrect costs. The more RI instances, the more exaggerated the error.

- **Customer Issue 110894**

Fixed: Turbonomic recommends invalid actions to provision new storage controllers. In most cases high utilization of storage controller compute resources does not warrant provisioning a new instance.

- **Customer Issue 110929,111051**

Fixed: For public cloud environments, under some circumstances when different pricing is available for a given workload instance type, Turbonomic can use the incorrect pricing to calculate the cost for associated RIs.

- **Customer Issue 110832,110848**

Fixed: Under some circumstances, when the topology component restarts, the user interface displays incomplete data, or even can be blank. The user interface should return with the last data that it showed before the need for a restart.

- **Customer Issue 110736**

Fixed: For OpenStack environments, under some circumstances discovery can fail to get CPU metrics, which triggers excessive logging.

- **Customer Issue 110775**

Fixed: For public cloud environments, under some circumstances after restarting the cost component, Turbonomic can fail to perform RI cost analysis. This can lead to not generating RI purchase actions when they are appropriate.

- **Customer Issue 110759**

Fixed: For Azure environments, under some circumstances with a large environment, timing issues cause analysis to use default policy settings even though custom policies are set for the given scope of entities.

- **Fixed:** For cloud environments, under some circumstances the Cloud Cost Comparison chart shows the results of actions that are disabled via policies.

- **Customer Issue 110849**

Fixed: For large environments, Turbonomic must improve performance when displaying the user interface.

- **Fixed:** For very large environments, Turbonomic must improve the performance of polling data to market analysis.

- **Customer Issue 110638**

Fixed: For Horizon VDI environments, if discovery fails for a single Business User, then discovery halts at that point, and it fails to discover the rest of the Business Users in the environment.

- **Customer Issue 110481**

Fixed: For NewRelic environments, under some circumstances discovery encounters entities managed by NewRelic that do not have a value for the entity Key. In that case, Turbonomic does not load the entity into the supply chain.

- **Customer Issue 110387**

Fixed: When you edit static groups, if the number of group members exceeds 20 (the count for paginating lists), then you cannot save those additional members to the group.

- **Customer Issue 110397**

Fixed: For deployments that include remote MySQL databases, database communications can generate a `BadSqlGrammarException` error.

- **Fixed:** For environments with Microsoft EA targets, under some circumstances EA discovery times out and fails to complete. The log includes messages similar to:

```
log.txt_20200415_174502:mediation-azureea-c8448796b-wz2r Caused by: java.net.SocketTimeoutException: Read timed out
```

If EA discovery fails to complete and you see this error, contact your support representative for assistance.

- **Customer Issue 110312**

Fixed: When you create an Optimize Cluster plan via the plan wizard, if you turn on **Ignore Constraints** the setting does not take effect. **Ignore Constraints** does work when you configure a custom plan.

- **Customer Issue 110014**

Fixed: For Dynatrace environments, under some circumstances Turbonomic displays fewer application entities than Dynatrace shows. This can occur when multiple services use the same process.

- **Customer Issue 109914,110889**

Fixed: For vCenter Server environments that support Auto Deploy, you can set up scenarios to patch your physical hosts with ESXi software. Under some circumstances, when the scenario brings the host back up from Maintenance mode, Turbonomic fails to recognize that the machine is back online.

Known Issues

- After you update Turbonomic from the 7.21 version family up to the 7.22 version family, when you review saved plans the plan results do not include Storage Amount data. To regenerate the Storage Amount data, run the plans again.

- The Shared-Nothing Migration feature is not fully supported at this time.

This release adds a new policy setting to turn on Shared-Nothing Migrations for a given scope of VMs. The What's New section of the User Guide describes this feature. However, this feature is not fully supported yet. In most cases you should not use the feature. If you want to experiment with Shared-Nothing Migration, please contact your support representative.

- Under certain circumstances, an update from a version earlier than 7.21.5 can cause schedules to be removed from scheduled policies.

This can happen if you created a schedule with the `z` timezone. This is not a valid timezone, and version 7.21.5 no longer supports it. After you update, you should check your scheduled policies to ensure the schedules are still attached.

- If you deploy Turbonomic to work with a remote database instead of the included database, then you must specify the correct SQL modes for the database. Configure the database to support:

```
STRICT_TRANS_TABLES,NO_ENGINE_SUBSTITUTION
```

In particular, the SQL modes should *not* include `ONLY_FULL_GROUP_BY`

- For Azure environments, Database resize actions do not properly consider storage capacity. As a result, Turbonomic can recommend resize down actions that are too aggressive, or it can fail to recommend appropriate resize actions. You should use `MANUAL` or `RECOMMEND` action modes, and verify that recommended actions are appropriate.

Turbonomic is aware of this problem and is working to address it as soon as possible.

- For Azure environments, under some circumstances you might need to manually start the analysis process to calculate Azure RI Buy actions.

Upon startup, because of timing issues for discovery of Azure billing entities (for example, Billing Families or costs), Azure RI analysis might not execute. Also note that if you restart the Turbonomic Action Orchestrator component for some reason, then Turbonomic can lose the currently pending Buy RI actions. Manually starting RI Buy analysis will correct these problems.

To manually force a restart of the RI Buy analysis, change the term of your Azure Profile settings. For example, navigate to **Settings: Billing and Costs: Reserved Instance Settings** and change the Azure Profile settings from **1 Year** to **3 Years**. This change triggers the RI Buy analysis.

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

Do not use the `numStorages` statistic filter. This statistic *will* be returned when other filters or no filters are used.

- The All Actions chart does not include pending actions for databases or database servers.
- For Azure environments with VMs in Scale Sets, for any VMs that are powered off the associated storage shows a utilization of zero GB. This is an accurate presentation of the data that the Azure environment returns for such a powered-off VM. However, it is likely that some of the storage capacity is currently utilized.
- **Customer Issue 110123**

There is a memory limit for the data you can download from the All Actions chart. For example, assume you have executed many actions over time in your environment. As a result, the list of all executed actions might exceed the data limit. In that case, downloading a CSV file from the All Actions chart will fail.

- Under rare circumstances the Turbonomic platform stops responding. This occurs when `etcd.service` fails. When it does occur, you should see the following error:

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

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

```
sudo systemctl restart docker.service
```

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

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

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

You should see messages similar to the following:

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

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

```
kubectl delete pod -n default gluster-{Unique_ID}
```

- When you update from 7.21.0 to this version, the update process sets your **Data Retention** setting back to the default values. If you have made custom data retention settings, you should reset them after you update.
- For public cloud environments, under some circumstances VMs or other workloads can appear in the user interface in an `UNKNOWN` state. However, there is no logging to indicate why the entity is in that state. This can occur when discovery succeeds, but the targeted infrastructure or cloud account cannot provide a known state for the workload. If this condition persists, you can contact your cloud vendor for more information about the workload.

- Updates from the 7.17 version family to the 7.21 version family cancel and delete any reservations that you have set up in the Placement view. If you require these reservations, then you should configure the reservations again in the updated version of Turbonomic
- If you are performing an *Online Update* from a 7.17 version of Turbonomic, then your update can inherit old configurations for the maximum MySQL connections. In large environments, or environments with many users of the same Turbonomic instance, this can result in the error, `error code [0]; Too many connections;`.

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

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

To correct this issue, change the configuration to allow 151 connections. The new setting should read `max_connections = 151`. After you make these changes, then restart the database. Either open a shell session as `root` or use `sudo`, and run the command, `systemctl restart mariadb`.

For assistance, contact Technical Support.

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

To resolve this problem, edit the All Actions charts to change them to Accepted Actions charts.

- Update Deletes Saved Plans

Because of changes to Turbonomic plans, when you update from the 7.17 version family to the 7.21. version family, the update process deletes your saved plans.

- For very large environments that use the WMI targets, the WMI discovery can run out of memory. To address this issue, Turbonomic has improved memory handling with WMI discovery.

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

- **Customer Issue 108841**

In NetApp environments, the storage controller shows 100% utilization when there are no more disks in a `SPARE` state that the storage controller can utilize in an aggregate. This does not indicate that the storage controller has no capacity.

- In vCenter Server environments, charts can show that a Virtual Datacenter (VDC) uses resources at more than 100% of capacity.

The utilization metrics that vCenter returns to Turbonomic for a VDC include utilization of resources that are reserved for vCenter overhead. However, the capacity metrics that Turbonomic discovers do not account for these reserved resources. As a result, it is possible Turbonomic shows that the VDC consumes more than 100% of capacity.

- In Azure environments, a subscription can use locked storage or locked resource groups. For such subscriptions, Turbonomic discovers incomplete data. Locked resources affect Turbonomic discovery in either of these scenarios:
 - A locked resource group

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

Turbonomic discovers all the entities in the resource group except the locked storage. It also discovers the resource group.

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

Typically, you will use templates that are discovered through your hypervisor targets. Along with discovering resource capacities for the given VM, Turbonomic should also discover the Image specification for a given discovered template. However in this version, Turbonomic does not discover the Image descriptions. In addition, discovered templates and their image specifications are read-only. For this reason, you cannot set up placement or reservations using discovered templates.

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

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

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

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

- Create a VM group that contains all the affected VMs. Note that Turbonomic automatically creates a group named `VMs_vCenter` that you might be able to use.
- Create a new VM automation policy. This policy will disable storage move actions.
- Set the group that you created to be the policy scope.
- Under **Action Automation** add the `Storage Move` action and set it to `Disabled`.
- In cases where actions recommend that you suspend hosts, the Optimal Improvements chart should indicate no utilization on the hosts to be suspended. Under some circumstances, the chart can show utilization on these hosts. The result is incorrectly low values for utilization on the other hosts in the current scope.
- Turbonomic generates special average or max utilization templates that it uses when calculating cluster headroom. You should not edit these templates, because Turbonomic will overwrite your changes the next time it generates the templates. However, the Template Catalog presents these templates as editable.

- After you run a plan, the user interface enables you to make changes to the plan configuration and then run the plan again. If you do this, the plan results will be inconsistent. If you want to run a plan with a different configuration, you should start a new plan.
- You should never use duplicate names for groups of the same entity type. However, the user interface does not validate group names to keep you from creating a duplicate name.
- For VMs running on Hyper-V, if you set a VCPU limit (limit VCPU to less than 100%), then the VCPU utilization data that VM returns to Turbonomic is not correct. As a result, Turbonomic will not recommend that you increase the VCPU limit.