



Turbonomic 6.0.1 Release Notes

October 31, 2017

This document describes issues that are addressed in Turbonomic 6.0.1 – Release Date: October 31, 2017. All builds are cumulative. Applying 6.0.1 onto any release of Turbonomic v5.0 or later will include all previous fixes. Please see the Turbonomic 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>

End of openSUSE Support

NOTE: As of Turbonomic version 6.0, Turbonomic no longer supports running the Turbonomic platform on openSUSE linux. Starting with this version, we only deliver Turbonomic VMs that run the CentOS operating system. If you need to upgrade from a version that runs on openSUSE, contact Technical Support for assistance.

Turbonomic ships as a pre-configured virtual machine image. For the last several years, that VM ran the openSUSE distribution of Linux. Starting with release 5.7, Turbonomic began a process to migrate away from openSUSE 13.2 and to base further releases on CentOS. You should have already migrated from openSUSE to CentOS before you attempt an upgrade to version Turbonomic version 6.0 or later.

For more information, and for a roadmap that describes this transition, see the Green Circle article, [Retirement of openSUSE Support](#).

Changeover of User Interface and API

As of Turbonomic version 6.0, we are changing the default interfaces to our new implementations, and retiring the older interfaces. These include:

- User Interface

The new user interface is implemented as HTML. This user interface is now the default – When you first log into a Turbonomic session, the browser now displays the HTML user interface.

You can still navigate to the legacy user interface. Click **SWITCH TO CLASSIC UI** at the bottom-left of the browser page to open the legacy user interface in a new browser tab. Note that we are on a roadmap to retire the legacy user interface. You should accustom yourself to the new user interface, and be sure to give us feedback.

- REST API

The new REST API returns data as JSON objects, and in many calls you pass data in via JSON objects. We use this API to support the new user interface – Anything you can do with the user interface should be possible directly via the API.

The new API uses Swagger to document its methods. To access the Swagger Doc, use the URL, `https://[Your_Turbonomic_IP]/vmturbo/apidoc/`. You can find additional API documentation in the Turbonomic 6.0 REST API Guide. Find this guide at the following location:

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

You can still use the old REST API. However, we will stop updating this XML-based API, and we are on a roadmap to retire it completely. You should begin converting any API scripts that you currently rely on so that they use the new API.

Configuration Requirements

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

Security Requirements for Browsers

For web browsers to communicate with Apache, the Apache configuration requires TLS versions 1.1 or later. To use versions of Microsoft Internet Explorer 9 and 10, you must enable TLS 1.1 or later (in **Internet Options > Advanced**). For more information, see the following Green Circle article:

[The DROWN Attack: Configuring your Turbonomic Web Security](#)

Updating the Tomcat Server

There are circumstances when you might choose to upgrade the Tomcat server on Turbonomic to a later version. In this case you must copy a local configuration file to the tomcat installation.

After you update the Tomcat server:

- Copy the file `/usr/libexec/tomcat/server.local` to `/usr/libexec/tomcat/server`
- To ensure that this server configuration file is executable, perform the command: `chmod 755 /usr/libexec/tomcat/server`

Storage Requirements for the Turbonomic Server

Turbonomic now states 150GB or greater as a requirement for disk storage. For Turbonomic servers hosted on VMware hypervisors, you should provide 150GB *plus* swap space to match the RAM allocation.

Transport Layer Security Requirements

Starting with version 5.4, 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 suddenly fails to validate after installing Turbonomic 5.4 or later, this is probably 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.

Enabling HTTP and HTTPS Proxies

Turbonomic supports the use of HTTP and HTTPS proxies for internet communication. However, you must edit the Tomcat Server configuration file to add the required system variables.

The file you must edit is on your Turbonomic server at `/usr/libexec/tomcat/server`

In this file, search for the `OPTIONS` statement. It should appear in the config file similar to the following:

```
FLAGS=" $( $JAVA_OPTS ) $CATALINA_OPTS "  
OPTIONS="-Dcatalina.base=$CATALINA_BASE ..."
```

Add the following flags to the `OPTIONS` statement, giving values for your proxies:

```
-Dhttp.proxyHost  
-Dhttp.proxyPort  
-Dhttps.proxyHost  
-Dhttps.proxyPort  
-Dhttp.proxyUser  
-Dhttp.proxyPassword
```

-Dhttps.proxyUser

-Dhttps.propyPassword

The resulting OPTIONS statement should be similar to the following:

```
OPTIONS="-Dcatalina.base=${CATALINA_BASE} \  
-Dcatalina.home=${CATALINA_HOME} \  
-Dhttp.proxyHost=111.10.10.123 -Dhttp.proxyPort=123 \  
-Dhttps.proxyHost=112.10.10.123 -Dhttps.proxyPort=456 \  
-Dhttp.proxyUser=user -Dhttp.proxyPassword=password \  
-Dhttps.proxyUser=user -Dhttps.propyPassword=password"
```

Note that the values you provide for this file must match the values you provide when specifying a proxy in the Turbonomic user interface. After you make these changes, restart the Tomcat server.

For further assistance, contact Technical Support.

Enabling HTTP to HTTPS Redirects for the REST API

Starting with version 5.9.1, Turbonomic redirects http requests to the user interface and the REST API over to HTTPS. To enable this feature, if you are updating from an earlier version of Turbonomic then you must restart the httpd service on the Turbonomic server.

Use the following command: `service httpd restart`

For more information, contact Technical Support.

Enabling Reports on Custom SELinux Platforms

When you install the provided Turbonomic OVA, the SELinux platform is configured to enable the generation and delivery of Turbonomic reports. However, if you install the Turbonomic software on a custom SELinux platform, you must perform configuration changes to enable reports.

To configure reporting, execute the following command with root privileges:

```
setsebool -P httpd_can_network_connect on
```

Further, on RHEL 7.1 you must change the context of the cgi-bin directory to enable the execution of cgi scrips by executing the following commands with root privileges:

```
semanage fcontext -a -t httpd_sys_script_exec_t /srv/www/cgi-bin/  
restorecon /srv/www/cgi-bin
```

If you need additional assistance, contact Turbonomic Technical Support.

Enabling IOPS and Network Monitoring for OpenStack Mitaka

The Target Configuration Guide gives instructions to connect to OpenStack targets. However, if you are running OpenStack Mitaka, you must perform additional configuration on the Mitaka platform to enable IOPS and Network data collection from Physical Machines. For those instructions, please see the Green Circle article, [Enabling OpenStack PM Metric Collection](#).

SMI-S Provider Versions for EMC VNX and EMC VMAX Storage Solutions

To connect to EMC VNX and VMAX disk arrays, Turbonomic uses EMC SMI-S providers that have the given disk arrays added to them. Note that VNX and VMAX support different versions of SMI-S Providers:

- VNX
For VNX and VNX2 arrays, use SMI-S version 4.6.2, based on Solutions Enabler 7.6.2. We have verified Turbonomic control of VNX block storage using SMI-S version 4.6.2 as a target.
- VMAX
For VMAX arrays, use SMI-S version 8.1, which is included in Solutions Enabler 8.1 – We have verified Turbonomic control of VMAX storage arrays using SMI-S version 8.1 as a target.

Update Recommendations

You can apply this update to any GA version of Turbonomic from version 5.9 or later, if it is running on CentOS.

NOTE: If your Turbonomic installation is running on an older version, or if it is running on openSUSE, contact Turbonomic Technical Support to confirm your update path.

Update Links

Turbonomic 6.0 is available as an offline update. Please see the Green Circle article: [How To Perform an Operations Manager "Offline Update" - Latest Links Included](#)

Related Green Circle Articles

- Offline and Online Update Instructions:
<https://greencircle.vmturbo.com/docs/DOC-1649>
- Release Notes and Product Documentation:
<https://greencircle.vmturbo.com/community/products/pages/documentation>

Known Issues

- To make placement decisions for workload on the public cloud, Turbonomic uses Cloud Budgets that you define. With these budgets, Turbonomic can assign costs and value to workload placement. To successfully run a Migrate to Cloud plan, you must first define a budget for the cloud environment that you will migrate to. Also, the budget value must exceed the costs for your current workload.
- To ensure that Turbonomic recommends the most appropriate actions for AWS environments, you must make specific settings for default policies in your Turbonomic installation. After adding an AWS target, then perform these steps:
 - Click **Settings: Policies** to navigate to the Policy Management page.
 - Ensure datastore browsing is disabled for the Storage Defaults policy.
Find the Storage Defaults policy and make sure that the option for Disable Datastore Browsing is ON.
 - Make the following settings for the Virtual Machine Defaults policy:
 - * For **Action Automation**, set the following actions to **Manual**:
 - * Resize Up
 - * Resize Down
 - * Storage Move
 - * Start
 - * For **Resize Recommendation Incrementing Constants**, drag the Rate of Resize slider to **High**.

In addition, you must ensure that the Cloud Budget for the AWS regions is set to a value three times the cost for the workloads in your AWS accounts.

- Turbonomic includes a new setting for Automation Policies on VMs that is named Resize Target Utilization. This sets the percentage of resource utilization that Turbonomic considers to be optimal for the VMs in the policy's scope. If you set this to a policy, then your change will affect analysis and action recommendations for the real-time environment. However, this setting will not affect analysis in plans.
- Turbonomic introduces a scaling constraint that you can use to exclude specific templates from move actions (VM, only) or resize actions (VMs and database instances on the cloud). For example, workload resizing in the cloud changes the template that defines the workload. To constrain resize actions for a scope of workloads, you can exclude certain templates for that scope.
When you set policies to exclude templates for resize actions, the template exclusion initially fails. If this occurs, you can force Turbonomic to recalculate actions, and those actions will then recognize your template settings. To do this, change the action mode for an action in this policy, and apply that change. Then reset the action mode to its original value and apply it again.
- When running plans to migrate workload to the cloud, be sure to choose the **Migrate to Cloud** option, and do not use the **Workload Migration** option. If you use **Workload Migration**, you can successfully set up a migration to the cloud, but the resulting plan might not choose the least expensive regions for workload placement.
- Under some circumstances, the user interface can show an AWS Relational Database using resources from a Host entity, and it can recommend actions to move Relational Databases to hosts. In the cloud environment, the database should use resources from a VM. If you see these symptoms, please contact Technical Support.
- In AWS environments, Turbonomic discovers data from a given AMI to determine the OS of an associated VM. The OS influences cost calculations.
If you delete the AMI from your environment, then Turbonomic cannot discover the OS, and so it assumes a free Linux OS. This can result in incorrect calculations of cost in plans and in evaluations of real-time costs.
- For a Migrate to Cloud plan, Turbonomic calculates VM moves to use the best-priced template, and then calculates the best storage that is available from the new VM location. In some cases, the resulting storage price might not be the lowest possible.

- For a Migrate to Cloud plan, you can select a group of workloads to migrate, or you can select all VMs in your environment and migrate them. You should not run a plan to globally migrate all of your discovered VMs to the cloud. For this release, such a plan can give unpredictable results. However, migrating a group of VMs to the cloud runs successfully.
- For this release, **Add Workload to Cloud** plans do not work properly. They cannot place the new workload on the cloud.
- To enable the best resize-down performance for VCPU on the cloud, you should set the VCPU Resize Increment to 1 MHz. In the HTML user interface, you can set this globally. To set the VCPU increment for a specific scope, use the legacy user interface.
- For Migrate to Cloud plans, when you migrate a VM that has an attached ISO image, the plan shows a move of a non-existent OGB disk. This OGB disk is a representation of the ISO image, and you can ignore the move action. The other plan actions for the VM are correct.
- Under some circumstances, for Migrate to Cloud plans that migrate to multiple public cloud accounts, the plan results can show more VMs than you chose to migrate. However, the plan costs and actions are correct for the number of VMs you migrated.
- For a Migrate Workload plan, the plan scope must include both the workloads you want to migrate, and also the providers that you want to migrate the workloads to. If you do not set the scope in this way, then the plan might not show any migration results.
- To set up a Migrate to Cloud plan, you specify the scope of VMs that you want to migrate to the cloud. Currently, the plan configuration only supports a scope of one group or one VM. If you want to migrate workloads that are currently in different groups, then you can create a static or a dynamic group to contain those workloads and migrate those workloads to the cloud.
- Turbonomic supports logging in to AWS targets via AWS Identity and Access Management (IAM) with IAM Users or IAM Roles. To enable using Roles, you must run the Turbonomic software on an EC2 instance in the AWS cloud, and you must have the Turbonomic instance run as the IAM Role, and connect to the AWS target accounts with the appropriate IAM Role. To perform these actions, please contact Technical Support.
- For AWS clouds, Turbonomic can get the information it needs to display license costs for database instances. However, for Azure clouds that information is not available. For that reason, Turbonomic does not display license costs for Azure database instances.
- After Turbonomic discovers the VMs in your environment, you can rename these VMs via the hosting hypervisor. In that case, because of caching issues the user interface can display the name incorrectly. If the VM names in the user interface appear to be wrong, try refreshing your browser cache.
- Full management of storage on the public cloud includes identifying wasted storage. To enable this, create a policy for your cloud storage. Then under Storage Settings, add Disable Datastore Browsing and then turn it off. Turbonomic discovers groups of storage by cloud provider or cloud region, to make it easy for you to create this policy. Set these groups to be the policy scope.
- Turbonomic automatically creates some groups of on-prem VMs that do not work in Migrate to Cloud plans. If you use these groups, the migration plan will not complete. Specifically, Turbonomic creates VM groups for each target. These group names begin with the prefix `VMs_VMTTarget_`. Do not use these groups as the source of workloads to migrate to the cloud in a migration plan.

- For Tomcat, SQLServer, WebSphere, and other application or database targets that use a scope to identify target instances, Turbonomic can fail to validate or discover the targets. If you add a target via scope, and that scope does not have any VMs to host the target applications, then the target will not validate. If you later add hosts for the applications to that scope, Turbonomic does not dynamically recognize the change and then validate and discover. Even if you execute a Validate command for that target, Turbonomic can validate but it will not run discovery.
To avoid this problem, make sure your applications are running on hosts before you configure the target. If you have encountered this problem (you added hosts to a scope after configuring the target), delete the target from Turbonomic and set a new target with this scope.
- The user interface includes a mode called TV Mode. This mode is still under construction, and you shouldn't use it.
- Upgrades from an earlier version can result in the database storing plans from the earlier version, yet the user interface does not show these plans. This is known when upgrading from 5.8.5 to 5.9.1. It might also occur for any upgrades with plans stored in the database.
- If Turbonomic discovers a WebLogic target, and then the WebLogic and Turbonomic instance are both shut down, then Turbonomic will not rediscovers the WebLogic target after it restarts. To correct this issue, execute a rediscovery of the WebLogic target.
- For UCS Director targets, the documentation does not describe a known configuration requirement. In UCS Director the configuration includes a number of Physical and Virtual Accounts for different services. For example, UCS Director can include Physical Accounts for storage controllers and Virtual Accounts for hypervisors. These accounts can identify the services by IP address or DNS name. To use this UCS Director target in Turbonomic, the same services should also be set up as Turbonomic targets. To correctly configure these targets, they must use an IP Address or DNS Name that matches the identification that is specified in UCS Director Physical or Virtual Account.
For example, if a storage controller's Physical Account uses a DNS name, then you must also use that DNS name when you configure the same storage controller as a Turbonomic target.
- In RedHat installations, generated PDFs (such as generated reports) might fail to load. The workaround is to change the ownership of the `/srv/reports/pdf_files/` directory from `apache` to `tomcat`. For assistance, contact Technical Support.
- In the New User Interface, you cannot add a VMM target unless you provide a value for Ignore Hosts. To add a VMM target that does not ignore any hosts, provide an invalid host address (an address that does not point to any existing host).
- For WebSphere, Tomcat, and SQLServer targets, you can set a scaling policy to be horizontal or vertical. To set scaling for these targets, you must set it both in the Turbonomic user interface, and also in the properties file for the given probe. This means that for any of these targets, all discovered entities must use the same scaling property. For help with scaling policies for these targets, contact Technical Support.
- Turbonomic has improved the implementation of target probes for Tomcat, Websphere, and SQLServer. One effect is that these new targets use scope to identify the application servers, and they do not use IP addresses. If you have already configured these targets using their IP addresses, then you must reconfigure these targets using Scope to identify them.
- Earlier versions of Turbonomic on the CentOS platform included an Apache configuration that enabled directory browsing. Version 5.9 fixes this issue. If you migrate to a new CentOS VM, the fix will take effect. However, if you update an older version to a version containing this fix, you must manually restart the httpd service to have the fix take effect.
- When using the xml-based Turbonomic REST API (prior to the HTML User Interface), be aware that requests include the username and password for your Turbonomic user account. You should always use the secure HTTPS protocol to ensure these requests do not include the password in clear text.

- Under some circumstances when using the Deploy View, the user interface can fail to respond. After you request a placement, if the placement recommendation does not appear within one to two minutes, reload the browser running Turbonomic to reset the user interface.
- If you ignore hyperthreading for a scope that is less than a datacenter and then restart tomcat for the Turbonomic application, then Turbonomic does not calculate the CPU capacity correctly for the affected scope of hosts. If this problem occurs, perform a full rediscovery to show correct CPU capacity.
- There is a rare case where Turbonomic can lose the cluster relationship for VMs running on a host. Assume you move a host out of a cluster, and directly into the datacenter (not into a cluster). Next you create a new cluster and then move the host into that cluster. In this case, the VMs on that host will not belong to any PM cluster. This can also affect Provider Virtual Datacenters that belong to the new cluster.
To avoid this problem, create the cluster first, and move the host directly into it. If this problem occurs, rediscover your environment to establish the correct cluster relationships.
- For IBM PowerVM targets, Turbonomic creates a group named AIX-vioses-auto to contain discovered VMs. Under some circumstances, the PowerVM configuration does not provide the information Turbonomic needs to create this group. To fix this issue, see the following IBM workaround: <http://www-01.ibm.com/support/docview.wss?uid=isg3T1024482>.
- For VMs running on Hyper-V, if you set a VCPU limit (limit VCPU to less than 100%), then the VCPU utilization data that VM returns to Turbonomic is not correct. As a result, Turbonomic will not recommend that you increase the VCPU limit.
- For Hyper-V and VMM targets, secure connections use port 5986 and unsecure connections use 5985. As you enable or disable a secure connection, the user interface does not automatically set the correct port number. You must be sure to specify the correct port for the type of connection you want.
- In OpenStack environments, it is possible to place a VM belonging to a specific cluster (a VM flavor that is set via extra specifications) onto a host that is not a member of that cluster. Turbonomic cannot identify this configuration error, and will not recommend a move to place the VM on an appropriate host.
You can implement placement policies to ensure that VMs always get placed on the correct hosts. This can work even if there is no host cluster to match the VM flavor. However, in that case the user interface will not show these VMs as members of a PM cluster.
- In the JSON REST API, you can execute a DELETE method for Market and Market_default. The API should prohibit DELETE methods for these resources. If you are programming with the API, do not execute a DELETE on these market resources.

