



# Turbonomic 6.0.14 Release Notes

June 1, 2018

This document describes issues that are addressed in Turbonomic 6.0.14 – Release Date: June 1, 2018. All builds are cumulative. Applying 6.0.14 onto any release of Turbonomic v5.9 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](mailto:support@turbonomic.com), or open a ticket at:

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

## What's New for Version 6.0

The 6.0 family of Turbonomic releases includes the following new features:

- User Interface Enhancements

The new HTML user interface is now the default. Thanks to customer feedback, we have improved a number of the views. New things to see include:

- Dashboards Page – Use this page to create custom dashboards that focus on specific areas of your environment.
- Reporting Page – Set up subscriptions to standard reports, and generate reports on-demand.
- Home Page Overviews – To help you focus on what's important to you, the Home Page now shows overviews for Hybrid, On-Prem, and Cloud environments. You can still use the Search Page to set a scope to your session and drill down to details.
- Policies and Settings – The new user interface now exposes a full range of settings for business rules and automation. Use these to set policies for different scopes in your environment.

- **Enhanced Public Cloud Support**

We continue to improve our support on the public cloud. With this release you can see:

- **Enhanced cost analysis**

When calculating placement and resource allocation for workloads on the cloud, Turbonomic uses real data from your cloud accounts to calculate actual cost. The information Turbonomic uses includes Compute Costs for the workload templates you use, License Costs for the workload OS, Storage Costs based on the storage tier you are using, and IP Cost.

For cloud storage, Turbonomic identifies the right storage tier for your workloads, ensuring you get storage performance without paying more than you need. It also discovers stranded, or "ghost" storage. When you delete a VM in your environment, you can forget to delete the attached storage. Turbonomic discovers and suspends this stranded storage so you don't have to keep paying for it.

In addition, when planning a migration from on-prem to the cloud, Turbonomic can identify stable on-prem workloads that are good candidates for Reserved Instance (RI) pricing. The resulting plan shows the cost savings you can expect by migrating these workloads to RI.

- **Relational Database Control**

Turbonomic discovers database instances in your cloud environment, and manages them as it does VM workloads. Turbonomic analysis determines correct scaling of your database services, provisioning new instances when necessary and recommending suspension of instances when appropriate.

- **Enhanced Target Support**

This version of Turbonomic introduces the following target enhancements:

- **EMC VPLEX** – VPLEX aggregates and refines data collected between connected Storage and Hypervisor targets. Turbonomic supports EMC VPLEX in a local configuration.
- **Cisco UCS Central** – Cisco UCS Central aggregates multiple Cisco Unified Computing System targets onto a single point of management.
- **Storage Logical Pools** – With Hitachi Data Systems and EMC VMAX 3, Turbonomic discovers and manages logical pools in your storage network.
- **Hyper-V** – This release adds support of Hyper-V 2016, and includes a number of improvements to our Hyper-V support.

- **Cloud-Native Enhancements**

For OpenShift and Kubernetes environments, Turbonomic imports Node and Pod affinity constraints, and uses those constraints in its analysis. In addition, Turbonomic creates groups for containers.

## 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 6.0.14 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>

# 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 version 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. For the user interface to display, this redirect must be in place.

If you are updating from a version of Turbonomic that is earlier than 5.9.1, then you must restart the httpd service on the Turbonomic server. A restart enables this redirect after such an update.

To restart the httpd service, use the following command: `service httpd restart`

For more information, contact 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.

## Fixed Issues

- **Fixed:** For environments that include Hyper-V and vCenter hypervisors, Turbonomic can discover one cluster through both hypervisors but assume they are different clusters. This occurs because of upper- and lower-case cluster names for the different hypervisors.
- **Fixed:** The HTML user interface does not show values for chassis power and cooling.
- **Fixed:** In large NetApp environments, Turbonomic posts an error when it tries to discover more than 500 entities. The NetApp API limits returned data for each request to 500 entities, and Turbonomic does not page such large requests.
- **Customer Issue 100627**  
**Fixed:** CPU Provisioned Used and Memory Provisioned Used settings made for a host do not take effect on the given host.
- **Fixed:** The API does not return necessary consumer data for requests to `.../entities/uuid`.
- **Fixed:** For very large environments, Turbonomic must improve management of queued discovery to ensure that it can discover all entities and metrics.

## Known Issues

- When you configure a plan that uses host templates to add hosts to the environment, you must ignore constraints. In the plan configuration, click **General** (in the plan configuration supply chain), and then turn on **Ignore Constraints**.  
The User Guide does not include this required step.
- Under some circumstances, after you set the action mode to Automatic for storage moves, the actions list shows associated actions in the Manual mode.
- In Microsoft VMM environments, you should not use VMM Dynamic Optimization to manage load balancing. The Dynamic Optimization actions can conflict with Turbonomic actions. If you must use Dynamic Optimization, you should set the aggressiveness to Low, and set the frequency to 60 minutes or longer.
- For public cloud environments, under some circumstances Turbonomic displays incorrect values in the Cloud Cost Comparison chart.  
To collect valid cloud cost statistics, the Hourly Saved Statistics must be set to a value greater than 24 hours. To check this setting, navigate to **Settings : Maintenance Options** in the user interface.  
It is also possible that after you make changes to RI Cost, the values in this chart can be incorrect. If you make changes to **Settings : Budget and Costs : RI COST**, then you should rediscover your cloud targets, and wait 24 hours for the cost values to align with the new settings.
- In AWS environments, under certain circumstances Turbonomic can show negative costs in charts for specific regions. This can occur when the AWS target indicates that it offers templates, but those templates are not truly available. In that case Turbonomic calculates negative values for costs in that region.
- The browser that you use for the Turbonomic user interface must be synchronized with the Turbonomic server to within one minute. Without this synchronization, Turbonomic can show incorrect metric values.

- When setting up a Migrate To Cloud plan, it is possible to choose a group of *cloud* VMs to migrate. You can run such a plan to simulate migration from one region to another. However, do not use the Migrate To Cloud plan to optimize workload placement for a complete service provider (do not choose from Providers instead of Regions or Groups when setting the plan scope). If you choose a cloud service provider for the scope of workloads to migrate, the plan can show inconsistent results.  
To plan for the best placement of workloads on an entire service provider, use the Optimize Cloud plan. When you configure such a plan scenario, be sure to remove all placement constraints.
- In OpenStack environments, when you set up reservations to deploy workloads via OpenStack templates you must constrain the deployment to the OpenStack datacenter.  
When you define the workload to deploy, you specify a template to deploy and any constraints that you want Turbonomic to respect. To deploy an OpenStack template, turn on the **Limit initial placement to locations you specify** constraint, and manually choose the datacenter or datacenters that support the OpenStack template.
- Turbonomic now requires HTTPS to connect to the user interface, and to connect to the API. Also, API connections no longer accept clear-text authentication, so you cannot include authentication in URLs to execute API commands. Instead, you should use curl commands to execute API commands.  
To support HTTPS, Turbonomic includes a self-signed certificate by default. We recommend that you install a certificate from a trusted Certificate Authority. If you do not install a trusted certificate, you can still use curl to execute API commands if you include the `-k` flag.
- After editing the IP address of a Turbonomic target or deleting a Turbonomic target, we recommend that you restart Tomcat in order for the Supply Chain to correctly reflect the changes.
- Turbonomic policies include a default policy named **Global Actions Mode**. You can use this policy to globally disable all actions in Turbonomic. If you turn on the setting to disable all actions, then `Disabled` takes effect for all actions. However, in the default policies the action modes remain as you have set them. The user interface does not show that you have globally disabled these actions.
- For cases where actions indicate provisioning new hosts, the Optimized Improvements chart does not include the hosts to provision in the After Plan section.
- If you use Google Chrome for the Turbonomic user interface, you must turn off the Chrome Preview mode before you download reports.
- In AWS environments that use RI templates across access regions, you should be sure to have a single master account, and include that master and all sub accounts as Turbonomic targets.  
If you experience situations where RI actions seem to recommend templates that are not available in the indicated regions, verify that you have included all the sub accounts as targets. If the problem persists, contact Technical Support.
- Under some circumstances in AWS environments, Turbonomic cannot recommend changing a workload from on-demand billing to an RI purchase plan. This occurs when the name, description, and platform fields of the workload's image do not clearly identify the OS. In that case, Turbonomic cannot identify the RI that matches the given workload.  
To enable RI optimization for all of your AWS workloads, make sure the image fields include a clear OS name such as `ubuntu`, `centos`, or `Windows 10`. If your account can find an RI that matches the identified OS, then Turbonomic can recommend switching that workload to an RI purchase plan.
- Do not make policy settings on parent groups (groups of groups).  
When you make policy settings on a parent group, Turbonomic resets the setting to the default every time you restart. This can have unintended consequences if you want these policies to implement business rules that control Turbonomic actions.  
When setting up policies, be sure you set them to individual groups. If necessary, create a custom group for the settings you want to apply.
- When using the Firefox browser, platform diagnostics cannot be downloaded. To download diagnostics, you must use another browser.

- The user interface cannot display Suspend actions for storage unless you make the following action mode settings for storage:
  - Delete set to Manual
  - Terminate set to Recommend
  - Suspend set to Recommend

After executing **Settings : Maintenance Options : Export Current Environment**, the user interface sometimes shows that the export failed even though the export is still running successfully.

If you are exporting a large topology, the user interface response can time out and show this warning after 60 seconds, even though the export continues to run.

If you experience this situation, restart the Turbonomic server. This restarts the HTTP server with a 10 minute timeout setting, which should be sufficient to export most environments. If the problem persists after a restart, contact Technical Support.

- Storage Suspend actions appear grouped with Delete Wasted Files actions in the Delete category of pending actions.

For example, assume you are viewing the On-Prem Overview on the Home Page, and the Pending Actions chart does not list Suspend Storage actions. However, if you hover on the Storage tier of the Supply Chain, the tooltip shows that you have actions to execute on Storage. These actions would be to suspend storage.

When viewing the overview, it is always a good idea to hover on any tier of the Supply Chain that is not completely green. The tooltip gives you extra information about that tier, including a count of actions.

To see the specific storage actions, click the Storage tier in the Supply Chain, and then view the Actions list.
- When running a Replace Workload plan to replace certain VMs with a template, under some circumstances the plan replaces the VM with a copy of an existing VM. For this reason, you should not configure a Replace Workload plan.

To get the same effect as a Replace Workload plan, you should configure the plan to remove the VMs you want to replace, and then configure it to add instances of a template or VM copy. For example, remove three VMs, and then add three instances of a VM template.
- For Load Balancer entities, the Transactions Per Second data can be incomplete.
- When setting up a custom plan, the user interface allows you to set a scope that includes public cloud entities. If you run such a custom plan, you can see unexpected or incorrect plan results. You should not set a scope that contains public cloud entities.

To run plans for the public cloud, choose the **Migrate to Public Cloud** or the **Optimize Cloud** plan type.
- In order to add an ACI Target, your Turbonomic instance must be using Market 1, which is disabled by default. If you require ACI Integration, please contact Turbonomic Technical Support.
- In AWS environments, a Reserved Instance (RI) can be associated with a region or an availability zone. For RIs associated with zones, the Turbonomic scale actions assume reservations are regional. As a result, actions may resize a VM based on its reserved cost, but the VM charges will actually use the on-demand cost.

If this occurs, use the AWS console to change the scope of the RI to "regional".
- 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 can show utilization on the remaining hosts as though the hosts to suspend are still active. The result is incorrectly low values for utilization on the hosts in the After Plan section of the chart.
  - When working with templates in the JSON API, the field names for some of the input DTO fields do not match the corresponding field names in the template output DTO.
  - Turbonomic can automate resize actions for datastores. However after executing the action, the hypervisor that is stitched to the datastore requires a refresh before it can recognize that change. Because Turbonomic uses the hypervisor to discover the datastore, then Turbonomic will not recognize the change either. As a result, Turbonomic might recommend the same action again, even though the datastore has already been resized. To avoid this situation, Turbonomic suggests that you set the Action Mode for storage resize actions to be `Recommend`.
  - When targets are added that introduce the Logical Pool entity to the Supply Chain, scoping to the Logical Pool entity will not show data for that entity.
  - In action scripts, you can get the internal name of a VM and use that to assemble calls to the API that work with the given VM. However, with the JSON API you must use the VM's UUID to access it via the API call to `entities`. Calls that use `$VMT_TARGET_NAME` with the XML API must now use `$VMT_TARGET_UUID`. For example, the following code creates the URL base for a call to the API that will get the actions associated with the action script's target VM:  

```
URL= "/vmturbo/rest/entities/$VMT_TARGET_UUID/actions"
```
  - 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.
  - **Customer Issue 98320**  
In Turbonomic you can create policies with a scope to a dynamic group. If changes in your environment remove all the members of the group for that policy, then Turbonomic deletes the policy definition.
  - 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.  
In Azure environments, when you set policies to exclude templates for resize actions on Database Instances, 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.
  - 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 enabled for the Storage Defaults policy.  
Find the Storage Defaults policy and make sure that the option for Disable Datastore Browsing is OFF.

- 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.
- 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.
- After restarting the Turbonomic server, users must log into new sessions in order to continue using the user interface or the API.
- To enable the best resize-down performance for VCPU on the cloud, you should set the global VCPU Resize Increment to 1 MHz.
- 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.
- 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.
- When you set up action orchestration in an Automation Policy, you should be sure that the scope is for this policy does not include conflicts with individual entities. If a single entity is in two scopes, and one scope enables orchestration while the other scope does not, then Turbonomic arbitrarily chooses either orchestration or a Turbonomic action mode.  
For example, assume two groups - GroupA and GroupB, and assume a host named MyHost is in both groups. If you configure action orchestration for hosts in GroupA but you do not for hosts in GroupB, then you cannot be sure that Turbonomic will call the action orchestrator for actions on MyHost.
- Assume you have application or database servers as targets, and they use dynamic groups to define their scopes for monitoring. If you add new application or database servers to these dynamic groups, or if you shut down and then restart an existing server, then Turbonomic fails to discover the change and these servers will not appear in the user interface. To resolve this problem, execute a manual rediscovery of the affected target.
- Under some circumstances Turbonomic can recommend different actions in real-time analysis and in plans. These differences can be for move actions in on-prem environments, and move or resize actions on the public cloud. If you see differences in these recommendations for plans vs. real-time analysis, check the VM Storage Move action modes and VM Host Move action modes for the affected scope. These action modes can be set to enable shared-nothing migration. Try disabling shared-nothing for on-prem, or enabling it for cloud. For more information about Shared-Nothing Migrations, see **Introducing Turbonomic: Turbonomic Actions : Shared-Nothing Migration Actions** in the User Guide.
- 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.
- 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.

- Assume you set up a scope that is less than a full datacenter. If you ignore hyperthreading in that scope 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.

- When you reserve resources for VMs that you will deploy in the future, Turbonomic initially calculates these reserved resources as CPU Provisioned, Memory Provisioned, and Storage Provisioned. However, these resources ultimately revert back to values that do not include the reserved VMs.
- 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 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.

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.

