



Turbonomic 7.21.0 Target Configuration Guide

Turbonomic, Inc

500 Boylston St, 7th floor
Boston, MA 02116 USA
Phone: (844) 438-8872
www.turbonomic.com

COPYRIGHT

Copyright © 2010 - 2020 Turbonomic, Inc. All rights reserved

END-USER LICENSE AGREEMENT

https://cdn.turbonomic.com/wp-content/uploads/Turbonomic_Click_Through_Customer-License.pdf

Contents

Target Configuration	5
Adding and Removing Target Virtual Management Servers.....	6
Hypervisor Targets	8
Hyper-V.....	12
Creating A Service User Account.....	15
vCenter Server.....	16
Creating A Service User Account In vCenter.....	20
Other Information Imported from vCenter.....	21
Cloud Management Targets	23
Private Cloud.....	23
OpenStack.....	25
Virtual Machine Manager.....	31
Storage Targets	36
EMC VMAX.....	38
HPE 3PAR.....	41
NetApp.....	45
Restricted Service Accounts In NetApp.....	48
Pure Storage.....	52
Guest OS Process Targets	55
AppDynamics.....	56
Dynatrace.....	60
SNMP.....	62
Enabling SNMP.....	65
WMI.....	66
Creating a WMI User Account.....	69
Enabling WMI.....	69
Database Server Targets	71
SQL Server.....	73
Appendix — Target Configuration	77
Cisco Unified Computing System.....	77
Enabling Windows Remote Management.....	78
Enabling WinRM Via Global Policy Objects.....	79
Enabling WinRM Via PowerShell.....	80
Secure Setup of WSMAN.....	81

Port Configuration.....	81
Sample OpenStack SNMP Configuration File for PM Metric Collection.....	82



Target Configuration

A target is a service that performs management in your virtual environment. Turbonomic uses targets to monitor workload and to execute actions in your environment. Target Configuration specifies the ports Turbonomic uses to connect with these services. You must install Turbonomic on a network that has access to the specific services you want to set up as targets.

For each target, Turbonomic communicates with the service via the management protocol that it exposes — The REST API, SMI-S, XML, or some other management transport. Turbonomic uses this communication to discover the managed entities, monitor resource utilization, and execute actions.

You can assign instances of the following technologies as Turbonomic targets:

- Cloud Managers
 - Microsoft System Center 2012 Virtual Machine Manager and System Center 2012 R2 Virtual Machine Manager
- Database Servers
 - Microsoft SQL Server 2008 R2, 2012, 2014, and 2016
- Fabric Managers
 - Cisco UCS 3.1+
 - HPE OneView 3.00.04+
- Guest OS Processes
 - AppDynamics 4.1+
 - DynaTrace 1.1+
 - SNMP
- Hypervisors
 - Microsoft Hyper-V 2008 R2, Hyper-V 2012, and Hyper-V 2012 R2
 - VMware vCenter 5.1, 5.5, 6.0, 6.5, and 6.7
- Orchestrator Targets
 - Action Script
- Storage Managers
 - NetApp Cmode/7mode using ONTAP 8.0+ (excluding AFF and SolidFire)
 - EMC VMAX using SMI-S 8.1+

- Pure Storage F-series and M-series arrays
- HPE 3PAR InForm OS 3.2.2+, 3PAR SMI-S, 3PAR WSAPI
- WMI Probe Supported Systems
 - Windows 2019
 - Windows 2016
 - Windows 2012 / 2012 R2
 - Windows 2008 R2
 - Windows 10
 - Windows 8 / 8.1
 - Windows 7

Transport Layer Security Requirements

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.

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 Turbonomic Technical Support.
- `Certificates do not conform to algorithm constraints`

To correct this error, refer to the documentation for your target technology 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.

Adding and Removing Target Virtual Management Servers

The target services your Turbonomic installation will manage appear in the Target Configuration list. You can add, remove, and edit entries in this list. Note that the target service's account must be configured with privileges that support the Turbonomic activities you want to perform. For example, the following list shows how vCenter privileges correspond to activities Turbonomic can perform:

- **Read Only** — Enables Turbonomic monitoring and simulation (what-if scenarios) only
- **vCenter Administrator** — Enables Turbonomic monitoring, simulation (what-if scenarios), and automation functions
- **Enable Datastore Browse** — Enabling this property for the account gives Turbonomic the privileges it needs to enable its storage management functionality

Adding Targets

To add a target service, click the **Target Configuration** button, provide the requested information, and click **Apply** to validate those targets and start a new discovery.

Typical information you provide includes:

- **Target Type** — Choose among the supported VM Management technologies (Hypervisor, Cloud Management, Load Balancer, etc.)

After you choose the technology, then choose the specific target type for that technology. For example, for Hypervisor technology, the types you can choose include vCenter, RHEV, Hyper-V, and XenServer.

- **Hostname or IP address** — The address of the target service you want to add
- **User Name** — A valid account username for the target service
- **Password** — A password for the target service account

Editing and Removing Targets

To edit a target entry, select it in the list and then click **Edit**. The **Target Configuration Form** opens, where you can make your settings.

To remove a target, select the entry in the list and then click **Delete**.



Hypervisor Targets

A hypervisor is a service that creates and runs virtual machines (VMs), providing the VMs compute and storage resources. When you connect Turbonomic to hypervisor targets in your environment, Turbonomic controls your environment, assuring application performance while also utilizing resources as efficiently as possible.

One of the first steps in any Turbonomic deployment is to connect to the hypervisors within your environment. Once connected, Turbonomic discovers the VMs, the physical machines that host the VMs, the datastores that provide storage resources to the physical machines, and the virtual datastores that provide storage resources to the VMs.

Supply Chain

Each hypervisor requires a physical machine (host) and one or more datastores to provide compute and storage resources. Virtual machines (VMs) run on those physical resources, and the VMs in turn provide resources to applications.

At the bottom of the supply chain, physical machines consume resources from data centers.

If your environment includes SAN technologies such as disk arrays, then the storage consumes resources from that underlying technology. If you add these storage targets, then Turbonomic extends the supply chain analysis into the components that make up the disk array. For more information, see [Storage Manager Targets \(on page 36\)](#).

Actions

Turbonomic recommends actions for the hypervisor supply chain as follows.

NOTE:

This is a general list of actions for entities discovered for hypervisors. Detailed actions per target are described in each target section.

Entity Type	Action
Virtual Machines	<ul style="list-style-type: none">• Provision additional resources (VMem, VCPU)• Move Virtual Machine• Move Virtual Machine Storage• Reconfigure Storage

Entity Type	Action
	<ul style="list-style-type: none"> Reconfigure Virtual Machine
Physical Machines	<ul style="list-style-type: none"> Start Physical Machine Provision Physical Machine Suspend Physical Machine
Storage	<ul style="list-style-type: none"> Start Storage Provision Storage Suspend Storage Move (only with Storage Targets configured) Resize (only with Storage Targets configured)
Consumer Virtual Datacenters	<ul style="list-style-type: none"> Resize Consumer vDC Provision Consumer vDC

Monitored Resources

Turbonomic monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz)

Entity Type	Commodity
	<ul style="list-style-type: none"> • IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) • Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) • Swap The utilization of the PM's swap space Measured in Kilobytes (KB) • Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) • CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Storage	<ul style="list-style-type: none"> • Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) • Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second • Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Datacenter	<p>NOTE: For datacenter entities, Turbonomic does not monitor resources directly from the datacenter, but from the physical machines in the datacenter.</p> <ul style="list-style-type: none"> • Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) • CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="362 233 1523 359"> <p>• IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="362 369 1523 495"> <p>• Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="362 506 1523 632"> <p>• Swap The utilization of the PM's swap space Measured in Kilobytes (KB)</p> <li data-bbox="362 642 1523 768"> <p>• Balloon The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB)</p> <li data-bbox="362 779 1523 947"> <p>• CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Kilobytes (KB)</p>
Provider Virtual Datacenter	<ul style="list-style-type: none"> <li data-bbox="362 974 1523 1100"> <p>• Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="362 1110 1523 1236"> <p>• CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="362 1247 1523 1373"> <p>• Storage The utilization of the storage attached to the Provider vDC. Measured in Kilobytes (KB)</p>
Consumer Virtual Datacenter	<ul style="list-style-type: none"> <li data-bbox="362 1404 1523 1530"> <p>• Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="362 1541 1523 1667"> <p>• CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="362 1677 1523 1803"> <p>• Storage The utilization of the storage attached to the Consumer vDC. Measured in Kilobytes (KB)</p>

Hyper-V

If you have a small number of Hyper-V hosts in your environment, you can add them individually as Turbonomic targets. Also, if you have deployed the Hyper-V hosts in a clustered domain (for example as a failover cluster), you can specify one Hyper-V host as a target and Turbonomic automatically add the other members of that cluster.

Note that for large Hyper-V environments, it's typical to manage the hosts via System Center Virtual Machine Manager (VMM). You can specify the VMM server as a target and Turbonomic will use it to discover and manage its child Hyper-V hosts. If you use VMM, you should not add individual Hyper-V hosts as targets. For information about adding VMM targets, see [Adding Virtual Machine Manager Targets \(on page 31\)](#).

NOTE:

Even if you manage Hyper-V using a VMM target, you must still configure remote management on each Hyper-V server. This Hyper-V topic includes instructions to configure remote management — see [Enabling Windows Remote Management \(on page 78\)](#).

For accurate SMB storage calculations, Turbonomic requires a VMM target.

Prerequisites

- Create a user account that Turbonomic can use to connect to your Hyper-V servers. See [Creating a Service User Account in Hyper-V \(on page 15\)](#)
- Configure remote management on each Hyper-V server. Refer to [Enabling Windows Remote Management \(on page 78\)](#)
- Your Hyper-V environment must not use Server Message Block (SMB) storage.
To manage SMB storage, Turbonomic requires a VMM target, and that VMM instance must manage the Hyper-V hypervisors and the SMB storage that they use.
Managing a Hyper-V plus SMB environment via Hyper-V targets will result in incorrect data collection for SMB storage.

Adding Hyper-V Targets

Once you've enabled remote management, you can add your Hyper-V hosts as targets. To add Hyper-V targets, select the **Hypervisors > Hyper-V** option on the Target Configuration page and provide the following information:

- Address
The FQDN of the Hyper-V host. If you're using the "Discover Host Cluster" below to add an entire cluster, enter the name of any one of the Hyper-V hosts in the cluster.
Note that you can enter an IP address for the host, but you must first configure an SPN on the host. Turbonomic recommends that you use the FQDN in this field.
- Port number
The port number for the remote management connection. The default HTTP port is 5985; the default HTTPS port is 5986.
- Secure connection
Select this option to use a secure connection (HTTPS). Make sure the required certificate is configured for use on the host.
- Full domain name

The full domain name of the cluster to which the host belongs.

- Discover Host Cluster

Turbonomic discovers and adds all Hyper-V hosts in the named cluster if this option is checked. Note that each server must be configured to allow remote management. You may find it helpful to configure WinRM using a GPO so new servers are configured automatically (see [Enabling WinRM Via a GPO \(on page 79\)](#)).

- Username

The username of a user account Turbonomic can use to connect to the Hyper-V host. If you checked "Discover Host Cluster" in the field above, use an account that is valid for all Hyper-V hosts in that cluster.

- Password

Password for account used.

Exporting Hyper-V Virtual Machines

In Hyper-V environments, you must be sure that all VMs have unique IDs.

Hyper-V supports the export of a VM, so that you can create exact copies of it by importing those exported files. The `Copy` import type creates a new unique ID for the imported VM. When importing VMs in your environment, you should always use the `Copy` import type.

Turbonomic uses the unique ID to discover and track a VM. If your environment includes multiple VMs with the same ID, then discovery will assume they are the same VM. As a result, the counts for VMs will be incorrect.

Supported Actions

For each discovered entity within the hypervisor supply chain, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations Only
Virtual Machine	Start, Move, Suspend, Resize Down, Resize Up	Terminate, Provision, Reconfigure
Physical Machine	Start, Suspend	Terminate, Provision
Storage		Provision

Monitored Resources

Turbonomic monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) Virtual Storage (VStorage)

Entity Type	Commodity
	<p>The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> Storage Access Operations Per Second (IOPS) <p>The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS</p> <ul style="list-style-type: none"> Latency <p>The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)</p>
Physical Machine	<ul style="list-style-type: none"> Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) Swap The utilization of the PM's swap space Measured in Kilobytes (KB)
Storage	<ul style="list-style-type: none"> Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second Latency The utilization of latency on the datastore

Entity Type	Commodity
	Measured in Milliseconds (ms)
Datacenter	<p>NOTE: For datacenter entities, Turbonomic does not monitor resources directly from the datacenter, but from the physical machines in the datacenter.</p> <ul style="list-style-type: none"> • Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) • CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) • IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) • Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) • Swap The utilization of the PM's swap space Measured in Kilobytes (KB) • Balloon The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) • CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Kilobytes (KB)

Creating A Service User Account

The service account Turbonomic uses to connect to a Hyper-V host must be an Active Directory domain account. The account must have full access to the cluster. To create such an account, execute the following command at a PowerShell prompt:

```
Grant-ClusterAccess <domain>\<service_account> -Full
```

Additionally, the service account must have specific local access rights on each host. The easiest way to grant Operations Manager the access it requires is to add the domain account to the `Local Administrators` group on each Hyper-V server.

Some enterprises require that the service account does not grant full administrator rights. In that case, you can create a restricted service account on every Hyper-V host:

NOTE:

Turbonomic does not support Restricted User Accounts on Windows 2012 Hyper-V nodes.

1. Add the service account to each of the following local groups:
 - WinRMRemoteWMIUsers__ (Or Remote Management Users)
 - Hyper-V Administrators
 - Performance Monitor Users

NOTE:

These groups are standard Windows Server 2012 security groups. If you are using an earlier version of Windows Server and do not see these groups, contact Turbonomic Support for assistance.

2. Grant permissions to the service account.

In the WMI Management console, grant the `Enable Account`, `Act as Operating System`, and `Remote Enable` permissions to the service account.

3. Configure the WinRM security descriptor to allow access by the service account:

- At a PowerShell prompt, execute `winrm configSDDL default`.
- In the "Permissions for Default" dialog box, grant the service account Read and Execute access.

vCenter Server

VMware vCenter Server provides a centralized management platform for VMware hypervisors. To manage your VMware environment with Turbonomic, you specify a vCenter Server instance as a target. Turbonomic discovers the associated infrastructure and performs intelligent workload management.

Prerequisites

- A user account Turbonomic can use to connect to your vCenter and execute actions (see [Creating a User Account In vCenter \(on page 20\)](#)).

General Considerations Before Target Addition

- **Linked vCenters:**

For linked vCenters, you must add each vCenter separately. Turbonomic communicates with each vCenter through a separate API endpoint, but aggregates the data from those vCenters and makes cross-vCenter actions possible.

- **Restricting Turbonomic Access to Specific Clusters:**

When you add a vCenter target, Turbonomic discovers all of the connected entities that are visible based on the account used to connect to the vCenter target. If there are clusters or other entities you want to exclude, you can do this by setting the role for the Turbonomic account to `No access` in the vSphere management client.

- **Shared Datastores:**

If you add more than one vCenter target that manages the same datastore, you may see conflicts in the metadata Turbonomic maintains for each vCenter. For example, datastore browsing may display a conflict between active and

wasted files, or each vCenter may define the same datastore as a member of a different storage cluster. Turbonomic recommends that you do not add multiple vCenter targets that manage the same datastore.

- **VSAN Permissions:**

In order to enable VSAN support and discover groups based on storage profiles, you must ensure that the user role Turbonomic is assigned has the `Profile-driven storage view` permission enabled. This permission is *disabled* in the built-in `readonly` role.

Adding vCenter Targets

To add vCenter targets, select the **Hypervisors > vCenter** option on the Target Configuration page and provide the following information:

- **Address**
The name or IP address of the vCenter server.
- **Username/Password**
Credentials for the user account Turbonomic can use to connect to the vCenter Server. Include the domain if required (`<domain>\<username>`).

vCenter Post-Addition Imports

In addition to discovering entities managed by the hypervisor, Turbonomic also imports a wide range of vSphere settings, such as Host and Storage DRS rules, annotations, Resource Pools, and DRS HA settings (See [Other Information Imported From vCenter \(on page 21\)](#)).

VMware vSphere 6.0 introduced the ability to move VMs between vCenters. If you enabled this feature in your VMware environment, you can configure Turbonomic to include cross vCenter vMotions in its recommendations.

To configure Turbonomic to support cross vCenter vMotion recommendations, you must create a Workload Placement Policy that merges the datacenters on the different vCenters, and then another policy to merge the given clusters. Also note that the merged clusters must use the same network names in the different datacenters. To create a Merge Policy:

1. In the Policy Management Tab, select **Placement Policy**.
2. For `policy type`, select **Merge**.
3. For `MERGE`, choose the merge type, and click **Select**.

To merge datacenters choose `Datacenter`, to merge Host clusters choose `Cluster`, and for storage choose `StorageCluster`.

4. Choose the specific datacenters or clusters to merge in this policy, then click **Select**.
5. Click **Save Policy**.

NOTE:

Since Turbonomic can only execute vMotions between clusters that use the same switch type (VSS or VDS), make sure any clusters you merge use the same switch type. Although Turbonomic will not initiate VSS → VDS vMotions, vSphere may do so. If this happens, Turbonomic displays a compliance violation notification.

Supported Actions

For each discovered entity within the hypervisor supply chain, Operations Manager can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations Only
Virtual Machine	Start, Move, Suspend, Storage Move, Resize Down, Resize Up	Terminate, Provision, Reconfigure
Physical Machine	Start, Suspend	Terminate, Provision
Storage		Provision

Monitored Resources

Turbonomic monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="362 235 1523 359"> • Swap The utilization of the PM's swap space Measured in Kilobytes (KB) <li data-bbox="362 369 1523 493"> • Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) <li data-bbox="362 504 1523 684"> • CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Storage	<ul style="list-style-type: none"> <li data-bbox="362 697 1523 821"> • Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) <li data-bbox="362 831 1523 955"> • Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) <li data-bbox="362 966 1523 1089"> • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second <li data-bbox="362 1100 1523 1224"> • Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Datacenter	<p data-bbox="383 1270 1523 1339">NOTE: For datacenter entities, Turbonomic does not monitor resources directly from the datacenter, but from the physical machines in the datacenter.</p> <ul style="list-style-type: none"> <li data-bbox="362 1350 1523 1474"> • Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) <li data-bbox="362 1484 1523 1608"> • CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) <li data-bbox="362 1619 1523 1743"> • IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) <li data-bbox="362 1753 1523 1877"> • Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="362 233 1521 359"> <p>• Swap The utilization of the PM's swap space Measured in Kilobytes (KB)</p> <li data-bbox="362 369 1521 495"> <p>• Balloon The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB)</p> <li data-bbox="362 506 1521 682"> <p>• CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Kilobytes (KB)</p>
<p>Provider Virtual Datacenter</p>	<ul style="list-style-type: none"> <li data-bbox="362 695 1521 821"> <p>• Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="362 831 1521 957"> <p>• CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="362 968 1521 1094"> <p>• Storage The utilization of the storage attached to the Provider vDC. Measured in Kilobytes (KB)</p>
<p>Consumer Virtual Datacenter</p>	<ul style="list-style-type: none"> <li data-bbox="362 1129 1521 1255"> <p>• Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="362 1266 1521 1392"> <p>• CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="362 1402 1521 1528"> <p>• Storage The utilization of the storage attached to the Consumer vDC. Measured in Kilobytes (KB)</p>

Creating A Service User Account In vCenter

The service account you use must have specific permissions on the vCenter. The easiest way to grant Turbonomic the access it requires is to grant full administrator rights.

Some enterprises require that the service account does not grant full administrator rights. In that case, you can create a restricted service account that grants the following permissions to enable the required Turbonomic activities:

Turbonomic Activity	Required Privileges
Monitoring	Read-only for all entity types
Recommend Actions	Read-only for all entity types
Wasted Storage Reporting	Datastore > Browse Datastore
Execute VM Move (vMotion)	Resource > Migrate Resource > Query Vmotion Resource > Modify Resource Pool Resource > Assign VM to Resource Pool
Execute VM Storage Move (svMotion)	Datastore > Allocate Space Datastore > Browse Datastore Resource > Assign VM to Resource Pool Resource > Migrate Resource > Modify Resource Pool Resource > Move Resource Pool Resource > Query VMotion Virtual Machine > Configuration > Change Resource Virtual Machine > Configuration > Swap File Placement
Execute VM Resize	Virtual Machine > Configuration > Change CPU Count Virtual Machine > Configuration > Change Resources Virtual Machine > Configuration > Memory Virtual Machine > Interaction > Reset Virtual Machine > Interaction > Power Off Virtual Machine > Interaction > Power On
Discover tags	Global > Global tags In addition, you must open ports 10443 and 7443 on the target server

Other Information Imported from vCenter

In addition to discovering entities managed by the vSphere hypervisors and their resources, Turbonomic:

- Imports any vSphere Host DRS rules when DRS is enabled, and displays them on the **Policy > Workload Placement** view under **Imported Placement Policies**. Imported rules are enabled by default, but you can disable specific rules if you want.

NOTE:

In vCenter environments, Turbonomic does not import DRS rules if DRS is disabled on the hypervisor. Further, if Turbonomic did import an enabled DRS rule, that somebody subsequently disables that DRS rule, then Turbonomic will discover that the rule was disabled and will remove the imported placement policy.

- Imports any custom annotations and displays related groupings in the **Inventory > Groups** tree view, under **VC Annotations**. The service account must enable the **Global > Global tag** privilege, and the target server must open ports 10443 and 7443.
- For vCenter Server versions 5.5 and later, discovers Virtual Machine Storage Profiles and displays them as groups anywhere that you can set scope. The groups appear under **VC Storage Profiles**. You can use these discovered storage profiles the same as any other groups — For example, to scope dashboards, or to set the scope for specific action policies.
- Discovers resource pools and displays them as folders in the Inventory tree and as components in the Supply Chain Navigator. If you have the Cloud Control Module license, Turbonomic manages resource pools as Virtual Datacenters (VDCs) and can recommend resize actions. Root resource pools appear as Provider VDCs in the supply chain, whereas child resource pools appear as Consumer VDCs.
- Imports vSphere HA cluster settings and translates them into CPU and memory utilization constraints. These are displayed as cluster-level overrides under **Folders** on the **Policy > Analysis > Host** view.



Cloud Management Targets

Cloud-based datacenters support scalability, resource pooling, multi-tenancy, and self-service management of virtual resources. Turbonomic supports the following cloud technologies:

- Microsoft Virtual Machine Manager (VMM)

Cloud infrastructures declare resource pools to serve business needs for multi-tenancy, authorized visibility into the infrastructure, and regional distribution of resources. These pools manage resources such as compute, storage, memory, and network bandwidth in ways that support the cloud's business requirements. Turbonomic discovers and manages these resource pools in relation to the whole datacenter. This enables Turbonomic to:

- Display cloud resources within each pool
- Analyze workload distribution across the entire datacenter
- Control the distribution of workload at each level of the cloud infrastructure

Private Cloud

A private cloud manages resources in pools to support multi-tenancy and self-service provisioning of virtual workloads. Turbonomic manages these resource pools in real time as demand fluctuates. This includes demand across resource pools, virtual datacenters (VDCs), and tenants.

On the private cloud, you can use Turbonomic to:

- Set up charge-back and show-back for private cloud or service-provider scenarios
- For service-providers, set up scoped views to limit exposure to the customer base
- Plan hardware requirements — the planning scenarios takes cloud architectures into account

Supply Chain

For private clouds, Turbonomic discovers resource partitions that are managed by the cloud manager, as well as the workload running on these partitions (the VMs and applications) and, where applicable, the supply that hosts

workload (the physical machines and storage). Turbonomic represents these partitions as the following types of Virtual Datacenters (VDCs):

- **Provider VDC**
 A collection of physical resources (PMs and datastores) within a private cloud. The cloud administrator has access to these resources, and defines the datacenter members. Administrators allocate Provider VDCs to manage resources that will be allocated to external customers through one or more Consumer VDCs.
- **Consumer VDC**
 A collection of resources that are available for customers to perform self-service management of workload through the cloud. It is an environment customers can use to store, deploy, and operate virtual systems. Consumer VDCs use the resources supplied by a Provider VDC.

Actions

Turbonomic recommends actions for private cloud infrastructures as follows:

Entity Type	Action
Provider Virtual Datacenters	Turbonomic does not recommend actions to perform on a Provider VDC. Instead, it recommends actions to perform on the devices running in the datacenter.
Consumer Virtual Datacenters	<ul style="list-style-type: none"> • Resize Consumer vDC • Provision Consumer vDC

Monitored Resources

Turbonomic monitors the following private cloud infrastructure resources:

Entity Type	Commodity
Provider Virtual Datacenter	<ul style="list-style-type: none"> • Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB) • CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz) • Storage The utilization of the storage attached to the Provider vDC. Measured in Kilobytes (KB)
Consumer Virtual Datacenter	<ul style="list-style-type: none"> • Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB) • CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz) • Storage

Entity Type	Commodity
	The utilization of the storage attached to the Consumer vDC. Measured in Kilobytes (KB)

OpenStack

To manage your OpenStack® environment, Turbonomic connects to the Keystone identity service endpoint. Through this connection, Turbonomic discovers the other services it needs to control your OpenStack environment. If it discovers the necessary services, then it considers the OpenStack target to be validated.

Turbonomic has been tested with hypervisors running OpenStack on Red Hat Enterprise Linux and Red Hat RDO.

Prerequisites

Turbonomic supports OpenStack version Icehouse or later.

You must have the following services installed and enabled, and Turbonomic must have access to each service at its administrative endpoint:

Service	Purpose
Keystone	Discovery of OpenStack tenants
Nova	Management of cloud computing systems
Cinder	Management of block storage
Ceilometer	Telemetry
Gnocchi	Telemetry- extended

To verify that these services are enabled, Open the Horizon dashboard to the Admin/System page, or execute the following command at the OpenStack command line:

```
[root@openstack ~:(keystone_admin)]# keystone service-list
```

You should see each of the required services in the list.

In addition, you must have the following meters enabled for Turbonomic to collect metrics for your OpenStack VMs. These meters are cumulative. For example, if you are running OpenStack Newton, you need the meters listed under Mitaka as well as those for Newton. For more information about each meter, please refer to the official OpenStack documentation:

Release	Meters
Mitaka or Earlier	<ul style="list-style-type: none"> • compute.node.cpu.percent • cpu_util • disk.device.usage • disk.device.read.requests.rate • disk.device.write.requests.rate • disk.read.bytes.rate • disk.write.bytes.rate

Release	Meters
	<ul style="list-style-type: none"> • disk.read.requests.rate • disk.write.requests.rate • hardware.cpu.util • hardware.disk.size.used • hardware.memory.buffer • hardware.memory.cached • hardware.memory.used • hardware.network.incoming.bytes.rate • hardware.network.outgoing.bytes.rate • hardware.system_stats.io.incoming.blocks.rate • hardware.system_stats.io.outgoing.blocks.rate • memory.resident • memory.usage • network.incoming.byte • ns.rate • network.outgoing.bytes.rate
Newton	<ul style="list-style-type: none"> • cpu_l3_cache • memory.bandwidth.total • memory.bandwidth.local • perf.cpu.cycles • perf.instructions • perf.cache.references • perf.cache.misses
Ocata	<ul style="list-style-type: none"> • network.incoming.packets.drop • network.outgoing.packets.drop • network.incoming.packets.error • network.outgoing.packets.error
Pike	<ul style="list-style-type: none"> • memory.swap.in • memory.swap.out
Queens	<ul style="list-style-type: none"> • disk.device.read.latency • network.outgoing.packets.drop • network.incoming.packets.error • network.outgoing.packets.error

Adding OpenStack Targets

To add OpenStack targets, select **Cloud Management > OpenStack** on the Target Configuration page, and provide the following information:

- Hostname or IP Address

Provide the public URL of the Keystone service. The default port is 5000 — Do not provide a port if you want to use the default. For the default port, validation assumes a standard HTTP connection. If you provide a port value (for a port other than 5000), validation assumes a secure HTTPS connection unless you specify the protocol.

For example, `10.10.123.45:5001` will use HTTPS, while `http://10.10.123.45:5775` will use HTTP.

- Username

The account must have an administrator role on the specified tenant. This account must be authenticated by OpenStack.

- Password

The password for the administrator account.

- Tenant Name

The organizational structure within the Compute service that you want to manage. In a basic OpenStack installation this tenant is usually named `admin`.

Collecting Physical Machine Metrics

To collect full metrics from physical machines in your environment and display them in Turbonomic, the following actions must be taken:

- `/etc/ceilometer/pipeline.yaml` must be modified to match the following snippets:

```

- name: meter_snmp
  interval: 600
  resources:
    - snmp://[Compute Node FQDN]
    - snmp://[Compute Node FQDN]
    - snmp://[Compute Node FQDN]

  meters:
    - "hardware.cpu*"
    - "hardware.memory.used"
    - "hardware.disk.size.used"
  sinks:
    - meter_sink
    - name: meter_snmp_rate
      interval: 600
      resources:
        - snmp://[Compute Node FQDN]
        - snmp://[Compute Node FQDN]
        - snmp://[Compute Node FQDN]
      meters:
        - "hardware.network*"
        - "hardware.system_stats*"
      sinks:
        - snmp_sink

- name: snmp_sink
  transformers:
    - name: "rate_of_change"
  parameters:
    source:
    map_from:
    
```

```

name: "hardware\\. (system_stats.io|network)\\. (incoming|outgoing)\\. (block
s|bytes)"
unit: "(blocks|B)"
target:
map_to:
name: "hardware.\\1.\\2.\\3.rate"
unit: "\\1/s"
type: "gauge"
publishers:
- notifier://
    
```

- After modifying this file, you will need to create a `snmpd.conf` file. For information about creating this file, see [Sample OpenStack SNMP Configuration File for PM Metric Collection \(on page 82\)](#).
- Upload the file you just created to `/etc/snmp`.
- Enable and start the SNMP service with the follow commands:
 - `systemctl enable snmpd.service`
 - `systemctl start snmpd.service`

After these modifications are made, OpenStack will collect all available metrics for physical machines.

Actions

Turbonomic recommends actions for the cloud target supply chain as follows.

Entity Type	Action
Virtual Machines	<ul style="list-style-type: none"> • Provision additional resources (VMem, VCPU) • Move Virtual Machine • Move Virtual Machine Storage • Reconfigure Storage • Reconfigure Virtual Machine
Physical Machines	<ul style="list-style-type: none"> • Start Physical Machine • Provision Physical Machine • Suspend Physical Machine
Consumer Virtual Datacenters	<ul style="list-style-type: none"> • Resize Consumer vDC • Provision Consumer vDC

Monitored Resources

Turbonomic monitors the following resources for the cloud target supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

Entity Type	Commodity
	<ul style="list-style-type: none"> • Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) • Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS • Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> • Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) • CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) • IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) • Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) • Swap The utilization of the PM's swap space Measured in Kilobytes (KB) • Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) • CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Datacenter	<p>NOTE: For datacenter entities, Turbonomic does not monitor resources directly from the datacenter, but from the physical machines in the datacenter.</p> <ul style="list-style-type: none"> • Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="362 233 1461 359"> <p>• CPU</p> <p>The utilization of the PM's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> <li data-bbox="362 369 1461 495"> <p>• IO</p> <p>The utilization of the PM's IO adapters</p> <p>Measured in Kilobytes per second (KB/s)</p> <li data-bbox="362 506 1461 632"> <p>• Net</p> <p>The utilization of data through the PM's network adapters</p> <p>Measured in Kilobytes per second (KB/s)</p> <li data-bbox="362 642 1461 768"> <p>• Swap</p> <p>The utilization of the PM's swap space</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="362 779 1461 905"> <p>• Balloon</p> <p>The utilization of shared of memory among VMs running on the host. ESX-only</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="362 915 1461 1083"> <p>• CPU Ready</p> <p>The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only</p> <p>Measured in Kilobytes (KB)</p>
<p>Provider Virtual Datacenter</p>	<ul style="list-style-type: none"> <li data-bbox="362 1115 1461 1241"> <p>• Memory (Mem)</p> <p>The utilization of the Datacenter's memory reserved or in use</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="362 1251 1461 1377"> <p>• CPU</p> <p>The utilization of the Datacenter's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> <li data-bbox="362 1388 1461 1514"> <p>• Storage</p> <p>The utilization of the storage attached to the Provider vDC.</p> <p>Measured in Kilobytes (KB)</p>
<p>Consumer Virtual Datacenter</p>	<ul style="list-style-type: none"> <li data-bbox="362 1545 1461 1671"> <p>• Memory (Mem)</p> <p>The utilization of the Datacenter's memory reserved or in use</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="362 1682 1461 1808"> <p>• CPU</p> <p>The utilization of the Datacenter's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> <li data-bbox="362 1818 1461 1904"> <p>• Storage</p> <p>The utilization of the storage attached to the Consumer vDC.</p>

Entity Type	Commodity
	Measured in Kilobytes (KB)

Virtual Machine Manager

In a VMM environment, the VMM management server processes commands and controls communications with the Hyper-V hosts. To manage VMM, you set the management server as a target. Turbonomic communicates with that target, and also with the Hyper-V hosts that the VMM server manages. For this reason, if you add a VMM target, you should not also add individual Hyper-V hosts as targets. You must grant Turbonomic access to the VMM management server, and also to all the associated Hyper-V machines.

Prerequisites

- VMM Dynamic Optimization disabled (recommended) or set to Low aggressiveness, with a frequency of at least 60 minutes.
- Configure remote management on the VMM management Hyper-V server. Refer to [Enabling Windows Remote Management \(on page 78\)](#)
- Apply necessary hot fixes on the VMM host

For example, if you are running VMM Server on a Windows Server version earlier than Windows Server 2012 R2, you must apply the hotfix referenced in the Microsoft Knowledge Base article #2842230 (<http://support.microsoft.com/kb/2842230>).

- Apply time synchronization

The VM that hosts Turbonomic must be synchronized with each target VMM management server. The Turbonomic Installation Guide includes instructions for synchronizing the clock on the Turbonomic server.

- PowerShell execution must be enabled on the VMM management server.
- Configure port access

WinRM uses ports 5985 and 5986 for standard and secure communications, respectively. The firewall on your VMM server must open these ports.

- Preexisting Hyper-V Targets

If any of the hosts that make up the VMM target were added separately as Hyper-V targets (and you do not plan to exclude these hosts), these targets must be deleted. Failure to do so will create duplicate entities in the market, which will negatively impact Turbonomic performance.

Adding VMM Targets

Turbonomic uses the address and credentials you provide to discover the VMM target. From the VMM target, Turbonomic gets the list of managed Hyper-V instances. It then uses that list to discover each Hyper-V instance. The Hyper-V credentials you provide must be valid for all of these machines.

Turbonomic will also import your Availability Sets, representing them as placement policies for the affected infrastructure.

To add VMM targets, select the **Cloud Management > VMM** option on the Target Configuration page, and provide the following information:

- The IP address or host name of the VMM management server
- Which port to use for the WSMAN connection

For a standard connection (HTTP) use 5985. For a secure connection (HTTPS) use 5986.

- Enable or disable a secure connection

If you enable a secure connection, then you must configure a certificate, and you must configure Turbonomic to communicate over HTTPS. For more information, see [Secure Setup for WSMAN \(on page 81\)](#).

Note that setting a secure connection for VMM does not also set secure connections for the underlying Hyper-V hosts. Any communications between Turbonomic and VMM will be secure. To configure secure connections to the underlying Hyper-V hosts, you must specify secure connections on each one.

- Full domain name for the user account

This domain name identifies the user account for Active Directory authentication.

- Login credentials for the Hyper-V servers that are managed by the VMM target (Optional)

Turbonomic must log into the Hyper-V servers that the VMM server manages. If you leave the Hyper-V credentials blank, then it will use the same credentials that VMM uses. If you provide Hyper-V credentials, then it will use that service account to log into every Hyper-V managed by the VMM.

Note that the service account Turbonomic uses to log into a Hyper-V host must satisfy certain requirements. For more information, see [Creating a Service User Account in Hyper-V \(on page 15\)](#).

- Hosts in the VMM target that you would like to exclude from Turbonomic.

Turbonomic will not discover or manage any hosts you enter in this field. These hosts do not participate in the market, and Turbonomic does not generate actions that involve these hosts. Enter multiple hosts in a comma-separated list. This field accepts both host names, and FQDNs (Fully Qualified Domain Names), and allows the use of the wildcard characters * and ?.

Microsoft VMM and Hyper-V VLANs

Windows Server Hyper-V provides support for VLANs on host and VM partitions. If your Hyper-V environment makes use of this VLAN support, then your VM moves must be sensitive to which hosts provide networking access to your defined VM networks. If a VM is a member of a given VM network, then any move of that VM must be to a host that has access to the same network.

For Hyper-V targets in a VMM environment, the Cloud Control Module is aware of the VM networks, and ensures that a move is to a host that provides connectivity over the given VM network. For example, if Host 1 and 2 provides connectivity to VM Network A, and Host 3 and 4 provide connectivity to VM Network B, Turbonomic will never recommend a VM with access to Network A, residing on Host 1, to move to Host 3. That would render the VM unable to communicate on Network A.

Configuring SMB 3.0 File Shares Discovery

With VMM, Turbonomic can discover SMB 3.0 shares as datastores, assuming these shares have been properly added to your VMM service center. When you add shares to your VMM environment, be sure to:

- Use the Fully Qualified Domain Name of the file server

As you add the associated file server to your VMM environment (via the Add Storage Devices Wizard), be sure to specify the FQDN of the file server on the Specify Discovery Scope page of the wizard. Do not use the file server's IP address.

- Ensure that file server names are unique

Do not specify file servers with the same name, even if they belong to different domains. Turbonomic requires the file server names to be unique.

For information about setting up SMB 3.0 shares, please refer to your Microsoft documentation. For example, see [“How to Assign SMB 3.0 File Shares to Hyper-V Hosts and Clusters in VMM”](#).

Actions

Turbonomic recommends actions for the cloud target supply chain as follows.

Entity Type	Action
Virtual Machines	<ul style="list-style-type: none"> Provision additional resources (VMem, VCPU) Move Virtual Machine Move Virtual Machine Storage Reconfigure Storage Reconfigure Virtual Machine
Physical Machines	<ul style="list-style-type: none"> Start Physical Machine Provision Physical Machine Suspend Physical Machine
Consumer Virtual Datacenters	<ul style="list-style-type: none"> Resize Consumer vDC Provision Consumer vDC

Monitored Resources

Turbonomic monitors the following resources for the cloud target supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> Memory (Mem) The utilization of the PM's memory reserved or in use

Entity Type	Commodity
	<p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> • CPU The utilization of the PM's CPU reserved or in use <p>Measured in Megahertz (MHz)</p> <ul style="list-style-type: none"> • IO The utilization of the PM's IO adapters <p>Measured in Kilobytes per second (KB/s)</p> <ul style="list-style-type: none"> • Net The utilization of data through the PM's network adapters <p>Measured in Kilobytes per second (KB/s)</p> <ul style="list-style-type: none"> • Swap The utilization of the PM's swap space <p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> • Balloon The utilization of shared memory among VMs running on the host. ESX-only <p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> • CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only <p>Measured in Megahertz (MHz)</p>
Datacenter	<p>NOTE: For datacenter entities, Turbonomic does not monitor resources directly from the datacenter, but from the physical machines in the datacenter.</p> <ul style="list-style-type: none"> • Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) • CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) • IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) • Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) • Swap The utilization of the PM's swap space

Entity Type	Commodity
	<p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> Balloon <p>The utilization of shared of memory among VMs running on the host. ESX-only</p> <p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> CPU Ready <p>The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only</p> <p>Measured in Kilobytes (KB)</p>
Provider Virtual Datacenter	<ul style="list-style-type: none"> Memory (Mem) <p>The utilization of the Datacenter's memory reserved or in use</p> <p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> CPU <p>The utilization of the Datacenter's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> <ul style="list-style-type: none"> Storage <p>The utilization of the storage attached to the Provider vDC.</p> <p>Measured in Kilobytes (KB)</p>
Consumer Virtual Datacenter	<ul style="list-style-type: none"> Memory (Mem) <p>The utilization of the Datacenter's memory reserved or in use</p> <p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> CPU <p>The utilization of the Datacenter's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> <ul style="list-style-type: none"> Storage <p>The utilization of the storage attached to the Consumer vDC.</p> <p>Measured in Kilobytes (KB)</p>



Storage Targets

Adding a storage Target enables Turbonomic to connect to your storage subsystem through an SMI-S provider API or through the controller's native API. Turbonomic uses the target's API to access information about each of the underlying disk arrays, and uses this information to set disk performance characteristics according to the type and capacity of storage. This leads to improved workload placement. Similarly, Turbonomic knows the relationships between storage controllers and disk arrays, and about the location of datastores within those arrays. This information also helps optimize workload placement.

The section below describes the storage supply chain. For information on how to add specific storage targets, the resources Turbonomic can monitor for the various supply chain entities, and the actions it can take to optimize the environment, refer to the target configuration instructions for your specific storage type.

Supply Chain

Storage targets (storage controllers) add Storage Controller and Disk Array entities to the supply chain. Disk Array entities in turn host Storage entities (datastores).

Entity Mapping

Turbonomic Mapping	EMC VMAX	HPE 3Par	NetApp	Pure
Storage	Volume (Regular, Thin, Meta)	Virtual Volume	Volume	Volume
Disk Array	Disk Group or Thin Pool	CPG	Aggregate	Shelf Array
Storage Controller	VMAX Array	Controller	Controller / Filer	Controller

Actions

Turbonomic recommends actions for storage targets as follows.

NOTE:

This is a general list of actions for storage managed by storage controllers. Specific actions Turbonomic can recommend, and which actions it can automate depends on the actual technology — Not all actions make sense for all types of storage. For example, Turbonomic can automate a datastore move across disk arrays or storage controllers for NetApp in C mode, but not for other storage technologies.

You can see how actions differ per technology in each section that describes adding a specific type of Storage Manager target.

Entity Type	Action
Storage	<ul style="list-style-type: none"> Start Storage Provision Storage Suspend Storage Move (only with Storage Targets configured) Resize (only with Storage Targets configured)
Disk Arrays	<ul style="list-style-type: none"> Provision Disk Array Start Disk Array Suspend Disk Array Move Disk Array (for NetApp Cluster-Mode, only) Move Virtual Machine Move Datastore
Storage Controller	<ul style="list-style-type: none"> Provision Storage Controller (recommendation only)

Monitored Resources

Turbonomic monitors the following storage resources:

Entity Type	Commodity
Storage	<ul style="list-style-type: none"> Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second Latency The utilization of latency on the datastore

Entity Type	Commodity
	Measured in Milliseconds (ms)
Disk Array	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> CPU The utilization of the Storage Controller's CPU in use Measured in Megahertz (MHz) Storage Amount The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller Measured in Megabytes (MB)

EMC VMAX

Turbonomic supports management of VMAX2 and 3 Series storage arrays. The VMAX series is a family of enterprise storage arrays designed for SAN environments. Turbonomic connects to VMAX storage systems via an EMC SMI-S provider that has the disk arrays added to it. A single SMI-S provider can communicate with one or more disk arrays. When you specify an SMI-S provider as a target, Turbonomic discovers all the added disk arrays.

Turbonomic will create Storage Groups based on the SLO levels defined in VMAX3 Targets. By default, Storage Move actions will respect these SLO levels based on the configured response time.

Prerequisites

- EMC SMI-S Provider V8.x

- A service account that Turbonomic can use to connect to the EMC SMI-S Provider (typically the default `admin` account)

Adding VMAX Targets

To add VMAX targets, select the **Storage > VMAX** option on the Target Configuration page and provide the following information:

- Address
The IP or host name of the SMI-S provider. If using HTTPS, you must follow the IP with `:PORT`.
- Username
The Username for the SMI-S provider.
- Password
The Password for the SMI-S provider.

Entity Comparison

After validating the new target, Turbonomic discovers the connected storage entities. This table compares terms used in EMC VMAX to those used in Turbonomic:

EMC VMAX Name	Turbonomic Entity
Volume (Regular, Thin, Meta)	Storage
Storage Resource Pool (VMAX3) / Thick Provisioned Pool (earlier)	Disk Array
Storage Group (VMAX3) / Thin Provisioned Pool (earlier)	Logical Pool
VMAX Array	Storage Controller

Supported Actions

For each discovered entity, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations only
Storage	Provision (Cloning), Delete	Resize (V-Volumes only)
Logical Pool		Resize
Disk Array		Provision

Monitored Resources

When calculating available storage, Turbonomic excludes disks devoted to the VMAX operating system by default. If these disks are assigned to new raid groups or storage pools, the capacity of those disks will then be considered when calculating the capacity of the Storage Controller.

Turbonomic monitors the following storage resources:

Entity Type	Commodity
Storage	<ul style="list-style-type: none"> • Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) • Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second • Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Logical Pool	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • Storage Amount The utilization of the logical pool's capacity. Measured in Megabytes (MB) • Storage Provisioned The utilization of the logical pool's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the logical pool. Measured in Operations per second • Latency The utilization of latency on the logical pool. Measured in milliseconds (ms)
Disk Array	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) • Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB)

Entity Type	Commodity
	<ul style="list-style-type: none"> Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	<ul style="list-style-type: none"> Storage Amount The utilization of the storage controller's capacity. Measured in Megabytes (MB)

HPE 3PAR

HPE 3PAR StoreServ systems use controller nodes to manage pools of storage resources and present a single storage system to consumers. Turbonomic communicates with the HPE 3PAR system via both the WSAPI and SMI-S providers that are installed on the 3PAR controller node.

Prerequisites

- SMI-S Provider enabled and configured on the controller node.
- WSPAI Provider enabled and configured on the controller node.
- A service account on the controller node that Turbonomic can use to connect to the SMI-S and WSPAI providers.

NOTE:

For discovery and monitoring, the Turbonomic service account must have the `Browse` permission on all monitored domains. To exclude domains from monitoring, the Turbonomic service account must have no permissions on those domains. For action execution, Turbonomic requires the `Edit` permission.

Setting Up the SMI-S Provider

The HPE 3PAR SMI-S Provider should be installed on the controller node. It is disabled by default — you must ensure that it is installed properly and running on the controller node.

To enable the SMI-S provider:

- Log into the HPE 3PAR Command Line Interface (CLI).
Open a secure shell session (ssh) on the controller node. Default credentials are `3paradm/3pardata`.
- Check the current status of the SMI-S provider.
In the shell session, execute the command, `showcim`.
- If the CIM service is not running, start it.
Execute the command `startcim` to enable the CIM service and the SMI-S provider.

To stop the SMI-S provider, execute the command `stopcim -f -x`.

Setting Up the WSAPI Provider

The HPE 3PAR WSAPI Provider should be installed on the controller node.

To enable the WSAPI provider:

1. Log into the HPE 3PAR Command Line Interface (CLI).
Open a secure shell session (ssh) on the controller node. Default credentials are `3paradm/3pardata`.
2. Check the current status of the WSAPI provider.
In the shell session, execute the command, `showwsapi`.
3. If the WSAPI service is not running, start it by executing the command `startwsapi`.
Execute the command `set wsapi -http enable` to allow only insecure connections, or `set wsapi -https enable` to allow only secure connections.

To stop the WSAPI provider, execute the command `stopwsapi -f`.

Adding HPE 3PAR Targets

To add an HPE 3PAR target, select the **Storage > HPE 3Par** option on the Target Configuration page and provide the following information:

- Address
The name or IP address of the 3PAR controller node.
By default, the controller provides SMI-S data over port 5988 (HTTP) or port 5989 (HTTPS). If your installation uses a different port for SMI-S, include the port number in the Address field.
- Username/Password
Credentials for a user account on the controller node.

After validating the new target, Turbonomic discovers the connected storage entities. This table compares terms used in HPE 3PAR to those used in Turbonomic:

HPE 3PAR Name	Turbonomic Entity
Virtual Volume	Storage
CPG	Disk Array
AO Configuration	Logical Pool
Controller	Storage Controller

Supply Chain

Storage targets (storage controllers) add Storage Controller, Logical Pool and Disk Array entities to the supply chain. Logical Pool and Disk Array entities then host Storage entities (datastores). For a visual representation, see the introductory [Storage Supply Chain \(on page 36\)](#).

3Par Adaptive Optimization

Adaptive Optimization (AO) for HPE 3Par enables management of data storage across two or three tiers. AO places storage regions on the appropriate tier in response to periodic analysis that AO performs.

To work with the storage in an AO group, Turbonomic:

- Discovers each Common Provisioning Group (CPG) as a disk array

In the Turbonomic user interface, these disk arrays do not host storage — They appear empty. Turbonomic will not recommend storage moves between these disk arrays, because such moves would conflict with AO block-level placement.

- Creates a single logical pool that hosts all the datastores in an AO group

This logical pool represents the AO group, and it includes all the member CPGs. Turbonomic considers this single logical pool when it performs analysis — It can recommend moving storage into or out of the AO group. Also, Turbonomic aggregates resource capacity in this logical pool. For example, the IOPS capacity for the AO logical pool is a combination of IOPS capacity for the constituent CPGs.

You can see the AO logical pool in the Turbonomic user interface. The display name for this logical pool is the name of the AO Configuration.

Supported Actions

For each discovered entity, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations only
Storage	Provision, Resize Up/Down, Delete	Move
Disk Array	Provision, Resize Up/Down, Delete	
Logical Pool		Provision, Resize Up/Down, Delete
Storage Controller		Provision

Monitored Resources

Turbonomic monitors the following storage resources:

Entity Type	Commodity
Storage	<ul style="list-style-type: none"> • Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) • Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second • Latency The utilization of latency on the datastore

Entity Type	Commodity
	<p>Measured in Milliseconds (ms)</p>
Disk Array	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) • Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second • Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Logical Pool	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • Storage Amount The utilization of the logical pool's capacity. Measured in Megabytes (MB) • Storage Provisioned The utilization of the logical pool's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the logical pool. Measured in Operations per second • Latency The utilization of latency on the logical pool. Measured in milliseconds (ms)
Storage Controller	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • CPU The utilization of the Storage Controller's CPU in use Measured in Megahertz (MHz) • Storage Amount

Entity Type	Commodity
	The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller Measured in Megabytes (MB)

NetApp

The Storage Control Module adds support for NetApp filers running the Data ONTAP operating system. NetApp storage controllers are Storage Virtual Machines that manage storage arrays — Vfilers for 7-Mode and Vservers for C-Mode. Turbonomic connects to these storage controllers to support NetApp targets in 7-Mode and Cluster-Mode (C-Mode).

Prerequisites

- Storage Control Module license
- Transport Layer Security (TLS) is enabled
- A service account Turbonomic can use to connect to the NetApp target

Enabling TLS

Starting with version 5.4, by default Turbonomic requires Transport Layer Security (TLS) version 1.2 to establish secure communications with targets. NetApp filers have TLS disabled by default, and the latest version they support is TLSv1. If your NetApp target fails to validate on 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 Turbonomic Technical Support.
- Certificates does not conform to algorithm constraints
To correct this error, refer to your 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.

For information about enabling TLS, see the Data ONTAP **System Administration Guide** for sections on the SSL protocol.

Service User Account — Administrator Role

To discover and fully manage NetApp disk arrays, Turbonomic must have a service account that grants privileges to execute commands through the NetApp filer's OnTap API (ontapi). In most cases, you can provide a user account with Administrator privileges:

- **NetApp 7-Mode:** Create the administrator account from the NetApp command line — For example:

```
useradmin user add Turbonomic -g Administrators
```
- **NetApp C-Mode:** Create the administrator account via the NetApp OnCommand System Manager, or from the NetApp command line — For example:

```
security login create -role admin -username Turbonomic -application ontapi -authmethod
password
```

If you prefer not to grant full administrator rights, see [Creating Restricted Service Accounts In NetApp \(on page 48\)](#)

Adding NetApp Targets

To add a NetApp target, select the **Storage > NetApp** option on the Target Configuration page and provide the following information:

- Address
The name or IP address of the NetApp Storage Controller.
7-Mode: Enter the storage controller address.
Cluster-Mode (C-Mode): Enter the cluster management address.
- Username/Password
Credentials for the NetApp service user account that you have configured for Turbonomic to use.

After validating the new target, Turbonomic discovers the connected storage entities. This table compares terms used in NetApp to those used in Turbonomic:

NetApp Name	Turbonomic Entity
Volume	Storage
Aggregate	Disk Array
Controller / Filer	Storage Controller

Supply Chain

Storage targets (storage controllers) add Storage Controller and Disk Array entities to the supply chain. Disk Array entities then host Storage entities (datastores). For a visual representation, see the introductory [Storage Supply Chain \(on page 36\)](#).

Supported Actions

For each discovered entity, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations only
Storage	Move (C-Mode only)	Move (7-Mode), Provision, Resize Up
Disk Array		Resize Up, Move (C-Mode only), Provision (C-Mode only)
Storage Controller		Provision

Note that for NetApp in C-Mode, Turbonomic can automate moving a datastore to a disk array on the same storage controller, as well as moves to a disk array on a different storage controller.

Monitored Resources

Turbonomic monitors the following storage resources:

Entity Type	Commodity
Storage	<ul style="list-style-type: none"> • Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) • Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second • Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) • Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) • Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second • Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> • CPU The utilization of the Storage Controller's CPU in use Measured in Megahertz (MHz) • Storage Amount

Entity Type	Commodity
	<p>The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller</p> <p>Measured in Megabytes (MB)</p>

Restricted Service Accounts In NetApp

While Turbonomic prefers a NetApp service account with administrator rights, it is possible to create an account that has limited access, by following the steps outlined below, depending on NetApp mode.

NetApp 9.x Restricted Service Account Setup

If you prefer to use a service account that does not have full administrator rights:

1. Log into the NetApp filer from a command shell.
2. Create a role and assign it permission to execute each of the following commands:

For example:

```
security login role create -role RoleName -cmddirname "storage aggregate show"
-vserver Cluster-Name
```

The required capabilities are listed below:

- cluster identity modify
- cluster identity show
- lun create
- lun igroup create
- lun igroup modify
- lun igroup show
- lun mapping create
- lun mapping delete
- lun mapping show
- lun modify
- lun show
- network interface create
- network interface delete
- network interface modify
- network interface show
- statistics show
- storage aggregate create
- storage aggregate modify
- storage aggregate show
- storage disk show
- system controller flash-cache show

- system node modify
 - system node show
 - version
 - volume create
 - volume modify
 - volume move modify
 - volume move show
 - volume move start
 - volume qtree create
 - volume qtree show
 - volume show
 - volume snapshot create
 - volume snapshot modify
 - volume snapshot show
 - vserver create
 - vserver fcp nodename
 - vserver iscsi nodename
 - vserver modify
 - vserver options
 - vserver show
3. For execution privileges, execute the following commands for the given role, where `Role-Name` is the name of the role you are creating, and `Cluster-Name` identifies the cluster you want the role to affect. You must execute these commands individually to set privileges that affect each individual cluster:
- `security login role create -role Role-Name -access all -cmddirname "volume offline" -vserver Cluster-Name`
 - `security login role create -role Role-Name -access all -cmddirname "volume unmount" -vserver Cluster-Name`
 - `security login role create -role Role-Name -access all -cmddirname "volume move" -vserver Cluster-Name`
 - `security login role create -role Role-Name -access all -cmddirname "volume delete" -vserver Cluster-Name`
4. Create a user that will use the newly-created role.
- For example:
- ```
security login create -User-Name Turbonomic -r TurbonomicRole
```
5. Enter a password for the new user when prompted.
6. Give the user access to the `ssh` and `ontapi` applications by using the following commands, replacing `Role-Name` and `User-Name` with the role and user you created:
- ```
security login create -role Role-Name -username User-Name -application ontapi
-authmethod password

security login create -role Role-Name -username User-Name -application ssh
-authmethod password
```

NetApp 7-Mode Restricted Service Account Setup

If you prefer to use a service account that does not have full administrator rights:

1. Log into the NetApp filer from a command shell.
2. Create a role with API privileges.

For example:

```
useradmin role add TurbonomicRole <capabilities>]
```

where <capabilities> is a comma-separated list of capabilities assigned to the role. The required capabilities are listed below:

- api-aggr-list-info
- api-disk-list-info
- api-fcp-node-get-name
- api-flash-device-list-info
- api-igroup-list-info
- api-iscsi-node-get-name
- api-lun-initiator-list-map-info
- api-lun-map-list-info
- api-lun-list-info
- api-net-ifconfig-get
- api-nfs-exportfs-list-rules-2
- api-options-list-info
- api-system-get-info
- api-system-get-version
- api-volume-list-info
- api-snapshot-list-info
- api-perf-object-get-instances
- api-perf-object-instance-list-info
- api-perf-object-counter-list-info
- api-qtree-list
- security-api-vfiler
- api-vfiler-list-info
- api-volume-options-list-info
- login-http-admin
- login-*

Note that the last login capability (login-*) may be necessary for external users.

Execution capabilities:

- api-volume-create
- api-volume-size
- api-volume-offline
- api-volume-online
- api-volume-destroy
- api-aggr-add

- `api-aggr-create`
 - `api-aggr-offline`
 - `api-aggr-online`
 - `api-aggr-destroy`
3. Create a group and assign the role.
For example:

```
useradmin add TurbonomicGroup -r TurbonomicRole
```
 4. Create a user that is a member of the group.
For example:

```
useradmin user add Turbonomic -g TurbonomicGroup
```
 5. Enter a password for the new user when prompted.

NetApp C-Mode Restricted Service Account Setup

If you prefer to use a service account that does not have full administrator rights:

1. Log into the NetApp filer from a command shell.
2. Create a role and assign it permission to execute each of the following commands:
 - `aggr-get-iter`
 - `igroup-get-iter`
 - `cluster-identity-get`
 - `lun-map-get-iter`
 - `net-interface-get-iter`
 - `storage-disk-get-iter`
 - `system-get-node-info-iter`
 - `volume-get-iter`
 - `vserver-get-iter`
 - `fc-node-get-name`
 - `flash-device-get-iter`
 - `iscsi-node-get-name`
 - `options-list-info`
 - `qtree-list-iter`
 - `system-get-version`
 - `lun-get-iter`
 - `snapshot-get-iter`
 - `perf-object-get-instances`
 - `volume-get-iter`
 - `volume-move-get-iter`
 - `volume-move-start`

For example, to enable volume offline, execute the following:

```
security login role create -role TurbonomicRole -access all -cmddirname "volume
offline" -vserver <cluster_name>
```

3. Create a user based on the role you create.

Give the user access to the `ssh` and `ontapi` applications. For example:

```
security login create -role TurbonomicRole -username Turbonomic -application ontapi -authmethod password
```

```
security login create -role VMTurboRole -username VMTurbo -application ssh -authmethod password
```

Pure Storage

Turbonomic supports management of Pure Storage FlashArray systems. Note that one Pure Storage target manages a single Pure Storage FlashArray instance. The storage devices in the array are all flash storage — to analyze IOPS capacity, Turbonomic uses the setting made for SSD Disks in the Policy View.

Because of the improved performance of Pure Storage arrays, Turbonomic intelligently moves more demanding workloads to these datastores. Turbonomic analysis is also able to incorporate Pure Storage deduplication and compression when recommending actions.

Prerequisites

- A service account Turbonomic can use to connect to the FlashArray

This account needs privileges to execute commands through the Pure Storage API — Typically the default `pureuser` administrative account.

Adding Pure Storage Targets

To add a Pure Storage target, select the **Storage > Pure Storage** option on the Target Configuration page and provide the following information:

- Address

The name or IP address of the Pure Storage FlashArray.

- Username/Password

Credentials for the service account Turbonomic can use to connect to the FlashArray. The Username must not contain the domain. For example, `Username=jjsmith` is correct, while `Username=myDomain\jjsmith` will result in a failure to validate.

- Secure connection

When checked, uses SSL to connect to the Pure target. Most Pure installations do not accept insecure connections. If you receive an error when adding the target with secure connections disabled, try re-adding with this option enabled.

After validating the new target, Turbonomic discovers the connected storage entities. This table compares terms used in Pure to those used in Turbonomic:

Pure Name	Turbonomic Entity
Volume	Storage
Shelf Array	Disk Array

Pure Name	Turbonomic Entity
Controller	Storage Controller

Supply Chain

Storage targets (storage controllers) add Storage Controller and Disk Array entities to the supply chain. Disk Array entities then host Storage entities (datastores). For a visual representation, see the introductory [Storage Supply Chain \(on page 36\)](#).

Supported Actions

For each discovered entity, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations only
Storage		Resize Up
Disk Array		
Storage Controller		Provision

Pure Storage assigns all the disks managed by a storage controller to a single array, with a fixed form-factor. There are no actions to perform for an array — For example, there is no action to move a disk array from one storage controller to another. Likewise, there are no actions to move or provision volumes because of the fixed form-factor.

Monitored Resources

Turbonomic monitors the following storage resources:

Entity Type	Commodity
Storage	<ul style="list-style-type: none"> Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> Storage Amount

Entity Type	Commodity
	<p>The utilization of the Disk Array's capacity. Measured in Megabytes (MB)</p> <ul style="list-style-type: none"> Storage Provisioned <p>The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB)</p> <ul style="list-style-type: none"> Storage Access Operations Per Second (IOPS) <p>The summation of the read and write access operations per second on the disk array Measured in Operations per second</p> <ul style="list-style-type: none"> Latency <p>The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)</p>
Storage Controller	<p>NOTE: Not all targets provide all possible commodities. For example, some storage controllers do not expose CPU activity. When a metric is not collected, its widget in the UI will display no data.</p> <ul style="list-style-type: none"> CPU <p>The utilization of the Storage Controller's CPU in use Measured in Megahertz (MHz)</p> <ul style="list-style-type: none"> Storage Amount <p>The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller Measured in Megabytes (MB)</p>



Guest OS Process Targets

Turbonomic supports the following Guest OS Process targets:

- SNMP (Simple Network Management Protocol)

Supply Chain

Guest Process targets add Application entities to the supply chain. You can navigate to each individual application server to see specific details.

Actions

Turbonomic recommends actions for the Guest OS Process supply chain as follows.

NOTE:

This is a general list of actions for entities discovered for applications. You can see how actions differ per technology in each section that describes adding a specific type of application target.

Entity Type	Action
Applications	<p>Without discovered Guest OS Processes or Application Servers, Turbonomic doesn't perform actions on applications. Instead, it performs resize actions on the host VMs. If host utilization is high enough on the physical machine running the application VM, Turbonomic may also recommend provisioning a new host.</p> <p>For specific application servers, see the individual application entry.</p>
Virtual Machines	<ul style="list-style-type: none">• Provision additional resources (VMem, VCPU)• Move Virtual Machine• Move Virtual Machine Storage• Reconfigure Storage• Reconfigure Virtual Machine

Monitored Resources

Turbonomic monitors the following resources for the Guest OS Process supply chain:

Entity Type	Commodity
WMI/SNMP Application	<ul style="list-style-type: none"> <li data-bbox="435 352 1271 485">• Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) <li data-bbox="435 491 1271 623">• Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="435 651 1271 783">• Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) <li data-bbox="435 789 1271 921">• Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) <li data-bbox="435 928 1271 1060">• Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) <li data-bbox="435 1066 1271 1199">• Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS <li data-bbox="435 1205 1271 1337">• Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

AppDynamics

Turbonomic supports workload management of the application infrastructure monitored by AppDynamics, via adding the AppDynamics instance to Turbonomic as a target.

The Turbonomic integration with AppDynamics provides a full-stack view of your environment, from application to physical hardware. With information obtained from AppDynamics, Turbonomic is able to make recommendations and take actions to both assure performance and drive efficiency with the full knowledge of the demands of each individual application.

In its default configuration, the AppDynamics target will collect up to 1100 AppDynamics nodes within the default collection period when a proxy is used, and up to 5000 nodes when no proxy is required. Larger AppDynamics environments are expected to take longer than one cycle to collect complete data.

Prerequisites

- SaaS or on-prem installation of AppDynamics, version 4.1+
- A valid AppDynamics user account.

For all types of application instances, the service account must have the `Read Only User` role. For monitoring database instances, this user must also have the `DB Monitoring User` role.

NOTE:

In newer versions of AppDynamics where these roles are available, they should be used instead:

- Applications and Dashboards Viewer
- Dashboards Viewer
- DB Monitoring User
- Server Monitoring role

AppDynamics Database Servers

AppDynamics also monitors database servers. In order for your database servers to be correctly stitched to the rest of your environment, you must:

- Enable enhanced metric collection.

For Hyper-V hosts, you must install Hyper-V Integration Services on the target VM hosting the database. For more information, please refer to the following integration services TechNet article:

<https://technet.microsoft.com/en-us/library/dn798297%28v=ws.11%29.aspx>

For VMware hosts, you must install VMware Tools on the target VMs.

- Ensure that the database name in AppDynamics is resolvable to an IP address by the Turbonomic instance.

You may need to make changes to your DNS or the file `/etc/resolv.conf` on the Turbonomic instance.

Entity Comparison

After validating the new target, Turbonomic discovers the connected applications. This table compares terms used in AppDynamics to those used in Turbonomic:

AppDynamics	Turbonomic
Business Application	Business Application
Node	Application Server / Application
Database	Database Server

Adding an AppDynamics Target

NOTE:

It is possible to monitor certain applications or database servers with both AppDynamics and Turbonomic, but this must be avoided as it will cause the entities to appear duplicated in the market.

If an application is monitored by AppDynamics, do not add it as a separate Turbonomic application target.

To add an AppDynamics instance as a target, specify:

- **Hostname or IP Address**
The host name or IP Address of the AppDynamics instance.
- **Username@Account**
Username and account ID with the necessary role(s). The format must be *Username@Account*.
- **Password**
Password for the account used to connect to the AppDynamics instance.
- **Port**
the port used to connect to the AppDynamics instance. By default, this is set to ports 80 (HTTP) and 443 (HTTPS).

NOTE: For SaaS-based AppDynamics instances, you must use port 443.
- **Secure Connection**
When checked, Turbonomic will connect via HTTPS. Make sure the required certificate is configured for use on the host.
- **Proxy Host**
The address of the proxy used for this target. Only fill out proxy information if you connect to the AppDynamics instance via a proxy.
- **Proxy Port**
The port to use with the proxy specified above. By default, this is 8080.
- **Proxy Username**
The username to use with the proxy specified above.
- **Proxy Password**
The password to use with the proxy specified above.

Actions

Turbonomic recommends actions for the AppDynamics supply chain as follows.

Entity Type	Action
Application / Application Server	<ul style="list-style-type: none"> • Resize Heap Recommendation only. • Resize Thread Pool Recommendation only. • Resize Connection Capacity Recommendation only. • Suspend VM Recommendation only. • Provision VM Recommendation only.
Virtual Machines	<ul style="list-style-type: none"> • Provision additional resources (VMem, VCPU) • Move Virtual Machine

Entity Type	Action
	<ul style="list-style-type: none"> • Move Virtual Machine Storage • Reconfigure Storage • Reconfigure Virtual Machine

Monitored Resources

NOTE:

The exact resources monitored will differ based on application type. This list includes all resources you may see.

Turbonomic monitors the following resources for the AppDynamics supply chain:

Entity Type	Commodity
Application	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem consumed from the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of VCPU consumed from the hosting VM Measured in Megahertz (MHz) • Transactions The utilization of the allocated transactions per second for the given application Measured in transactions per second • Heap The utilization of the application server's heap Measured in Kilobytes (KB) • Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) • Threads The utilization of the server's thread capacity Measured in Threads • Connection The utilization of the connection capacity. Only applicable to database servers Measured in Connections • Garbage Collection The percentage of server uptime spent garbage collecting Measured in of uptime (%)
Virtual Machine	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM

Entity Type	Commodity
	<p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> Virtual CPU (VCPU) <p>The utilization of the VCPU allocated to the hosting VM</p> <p>Measured in Megahertz (MHz)</p> <ul style="list-style-type: none"> Virtual Storage (VStorage) <p>The utilization of the virtual storage capacity allocated for the VM</p> <p>Measured in Kilobytes (KB)</p> <ul style="list-style-type: none"> Storage Access Operations Per Second (IOPS) <p>The utilization of IOPS allocated for the VStorage on the VM</p> <p>Measured in IOPS</p> <ul style="list-style-type: none"> Latency <p>The utilization of latency allocated for the VStorage on the VM</p> <p>Measured in milliseconds (ms)</p>

Dynatrace

Turbonomic supports discovery of applications that are managed by the Dynatrace platform. Turbonomic includes the discovered information about these applications in its calculations for VM actions.

Prerequisites

- A Dynatrace Server instance running in your environment.
This instance must be configured to monitor applications running in your environment.
- Managed VMs that are running applications managed by Dynatrace.
For Turbonomic to discover applications through Dynatrace, the applications must be running on VMs in your environment. Also, these VMs must be managed by hypervisors that you have configured as Turbonomic targets.
- A valid Dynatrace user account.
The account must provide read-only access to the application data that is gathered by the Dynatrace server.

Adding a Dynatrace Target

NOTE:

It is possible to monitor certain applications or database servers with both Dynatrace and Turbonomic. You should avoid such a configuration because it can cause Turbonomic to generate duplicate entities in the market.

If you monitor an application via a Dynatrace server, and you configure that Dynatrace server as a Turbonomic target, then be sure you have not added that application as a separate application target in Turbonomic.

To add a Dynatrace server instance as a target, specify:

- Hostname or IP Address

The host name or IP and endpoint, separated by a slash. For example, 10.10.10.10/e/b70e3eb2-e82b-4c13-a5a4-560d9865841r

- API Token

The token that Turbonomic will use to authenticate its calls to the Dynatrace API. This token must have permission to execute GET methods via the Dynatrace API.

- Proxy Host

The address of the proxy used for this target. Only fill out proxy information if you connect to the Dynatrace server via a proxy.

- Proxy Port

The port to use with the proxy specified above. By default, this is 8080.

- Proxy Username

The username for the account to log into the proxy specified above.

- Proxy Password

The password to use with the proxy specified above.

Actions

Turbonomic does not recommend actions for Dynatrace applications. However, it considers resource utilization by these applications when recommending actions for the underlying VMs.

Monitored Resources

NOTE:

The exact resources monitored will differ based on application type. This list includes all resources you may see.

Turbonomic monitors the following resources for the Dynatrace supply chain:

Entity Type	Commodity
Application	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) • Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second • Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) • Threads The utilization of the server's thread capacity

Entity Type	Commodity
	<p>Measured in Threads</p> <ul style="list-style-type: none"> Connection <p>The utilization of the connection capacity. Only applicable to database servers</p> <p>Measured in Connections</p> <ul style="list-style-type: none"> Garbage Collection <p>The percentage of server uptime spent garbage collecting</p> <p>Measured in of uptime (%)</p>
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory (VMem) <p>The utilization of the VMem allocated to the hosting VM</p> <p>Measured in Kilobytes (KB)</p> Virtual CPU (VCPU) <p>The utilization of the VCPU allocated to the hosting VM</p> <p>Measured in Megahertz (MHz)</p> Virtual Storage (VStorage) <p>The utilization of the virtual storage capacity allocated for the VM</p> <p>Measured in Kilobytes (KB)</p> Storage Access Operations Per Second (IOPS) <p>The utilization of IOPS allocated for the VStorage on the VM</p> <p>Measured in IOPS</p> Latency <p>The utilization of latency allocated for the VStorage on the VM</p> <p>Measured in milliseconds (ms)</p>

SNMP

Turbonomic will discover application and operating system resources using SNMP (Simple Network Management Protocol) on Linux hosts in your entire environment, or a portion of your environment, based on scope.

Prerequisites

- Underlying VM host targets added to Turbonomic (for discovery)
- SNMP service enabled/configured on target VMs

NOTE:

For Hyper-V hosts, you must install Hyper-V Integration Services on the target VMs. For more information, please refer to the following integration services TechNet article:

<https://technet.microsoft.com/en-us/library/dn798297%28v=ws.11%29.aspx>

For VMware hosts, you must install VMware Tools on the target VMs.

Adding SNMP Targets

To add SNMP targets, select the **Guest OS Processes > SNMP** option on the Target Configuration page and provide the following information:

- Target Name

The display name that will be used to identify the target in the Target List. This is for display in the UI only; it does not need to match any internal name.

- Scope

Turbonomic will search for SNMP applications using the port entered above on the VMs found within the set scope. This scope can be set to the entire environment, single or multiple clusters, or particular virtual machines. Select this option and choose the scope for application discovery.

NOTE:

A single scope (target) must not contain more than 500 virtual machines. Turbonomic recommends utilization of multiple WMI targets for environments exceeding this limit.

- Community

The SNMP v2c community string Turbonomic will use to connect to the application

- Port number

The Port number Turbonomic will use to connect to the virtual machines hosting the application

- Full Validation

When selected, Turbonomic will require every VM in the selected scope to be a valid target. If Turbonomic is unable to authenticate a VM in the scope, the target will not validate and data will not be collected.

- Enable SNMPv3

When checked, Turbonomic will use SNMPv3 to connect to the virtual machines in the selected scope

- SNMPv3 Username/Securityname

The Username or Securityname that Turbonomic will use to connect to the virtual machines hosting the application

- SNMPv3 Enable Privacy

When checked, Turbonomic will encrypt using the privacy password to keep the connection private.

- SNMPv3 Authentication Password

The Authentication Password Turbonomic will use to connect to the virtual machines in the selected scope. This allows requests to be authenticated, confirming the sender's identity.

- SNMPv3 Privacy Password

The Privacy Password Turbonomic will use to connect to the virtual machines in the selected scope. For requests to be encrypted, the SNMP manager and the SNMP agent must share knowledge of the privacy password associated with the username.

- SHA-1 For Authentication

When checked, Turbonomic will use SHA-1 Authentication to connect to the virtual machines in the selected scope. When unchecked, Turbonomic uses MD5

- AES For Privacy

When checked, Turbonomic will use AES 128 to connect to the virtual machines in the selected scope. When unchecked, Turbonomic uses DES

NOTE:

Turbonomic does not support AES 192 or 256 for SNMP targets.

Supported Actions

For each discovered entity within the application supply chain, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Action
Applications	<p>Without discovered Guest OS Processes or Application Servers, Turbonomic doesn't perform actions on applications. Instead, it performs resize actions on the host VMs. If host utilization is high enough on the physical machine running the application VM, Turbonomic may also recommend provisioning a new host.</p> <p>For specific application servers, see the individual application entry.</p>
Virtual Machines	<ul style="list-style-type: none"> • Provision additional resources (VMem, VCPU) • Move Virtual Machine • Move Virtual Machine Storage • Reconfigure Storage • Reconfigure Virtual Machine

Monitored Resources

Turbonomic monitors the following resources for the application supply chain:

Entity Type	Commodity
WMI/SNMP Application	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)
Virtual Machine	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="435 226 1263 357"> • Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) <li data-bbox="435 367 1263 497"> • Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS <li data-bbox="435 508 1263 638"> • Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

Enabling SNMP

Turbonomic requires that SNMP is enabled and configured in order to discover Guest Processes. While these steps will change slightly between OS versions, this topic gives you the general instructions:

Enabling SNMP v2

1. Obtain and install the SNMP server package for your version of Linux. This is commonly called `net-snmp` in most package managers. Also install the `net-snmp-utils` package if it is available
2. Configure the SNMP daemon by editing `/etc/snmp/snmpd.conf` or running `snmpconf -i` at the command line in order to verify that:
 - SNMPv2c is enabled
 - A read-only community name has been set. This community name will be used by Turbonomic to communicate with the VM

NOTE:

The community name you select must have at least read-only access to all OIDs in the system.

3. Configure the daemon to listen on a public interface. Most default installations only listen on `127.0.0.1`
4. Start the SNMP daemon process

These long form and short form sample `snmp.conf` files illustrate the necessary changes.

Long form, VACM:

```
# First, map the community name "mycommunity1" into a "security name"
# sec.name source community
com2sec notConfigUser default mycommunity1
# Second, map the security name into a group name:
# groupName securityModel securityName
group notConfigGroup v2c notConfigUser
# Finally, grant the group read-only access to any OID.
# group context sec.model sec.level prefix read write notif
access notConfigGroup "" any noauth exact all none none
```

Short form, older syntax:

```
rocommunity mycommunity1 default system
```

Verify Your SNMP Setup

Verify that your SNMP setup is successful by using the following command from a remote machine, replacing the community string and IP address:

```
snmpwalk -Os -v 2c -c COMMUNITY_STRING IP_ADDRESS iso.3.6.1.2.1.1.1
```

If successful, the command will return the kernel version of the machine (similar to the output of `uname -a`)

Enabling SNMP v3

1. Obtain and install the SNMP server package for your version of Linux.

Most package managers call this package `net-snmp`. To verify your configuration in step 5 below, you should also install the `net-snmp-utils` package on the Turbonomic VM.

2. Stop the SNMP service.
3. Create the SNMPv3 user.

Execute the following command:

```
net-snmp-config --create-snmpv3-user [-ro] [-A authpass] [-X privpass] [-a MD5|SHA]
[-x DES|AES] [username]
```

For example, the command:

```
i. net-snmp-create-v3-user -ro -A snmpv3authPass -a SHA -X snmpv3encPass -x AES
snmpv3user
```

results in output similar to:

```
adding the following line to /var/lib/net-snmp/snmpd.conf:
```

```
createUser snmpv3user SHA "snmpv3authPass" AES snmpv3encPass adding the following
line to /etc/snmp/snmpd.conf:
```

```
rouser snmpv3user
```

4. Secure the SNMP daemon.

Use a host firewall to only allow requests from source IP addresses you know, and which need to query the system.

5. Verify your SNMP setup.

To verify the configuration, you can execute the following command on the Turbonomic VM:

```
snmpwalk -v3 -u snmpv3user -A snmpv3authPass -a SHA -X snmpv3encPass -x AES -l
authPriv <TargetIP>
```

WMI

Turbonomic will discover application and operating system resources using WMI (Windows Management Instrumentation) in your entire environment, or a portion of your environment, based on scope.

Prerequisites

- Underlying VM host targets added to Turbonomic (for discovery)

- A WMI user account Turbonomic can use to connect to the WMI targets. This account can either be an administrator user, or a non-administrator belonging to certain local user groups. For a full list of requirements, see [Creating a WMI User Account \(on page 69\)](#)
- WMI enabled on target VMs. For assistance in enabling WMI, see [Enabling WMI \(on page 69\)](#).

NOTE:

For Hyper-V hosts, you must install Hyper-V Integration Services on the target VMs. For more information, please refer to the following integration services TechNet article:

<https://technet.microsoft.com/en-us/library/dn798297%28v=ws.11%29.aspx>

For VMware hosts, you must install VMware Tools on the target VMs.

Adding WMI Targets

To add WMI targets, select the **Guest OS Processes > WMI** option on the Target Configuration page and provide the following information:

- Target Name

The display name that will be used to identify the target in the Target List. This is for display in the UI only; it does not need to match any internal name.
- Username

The username Turbonomic will use to connect to the WinRM service on the Virtual Machine hosting the application. This should not include the Active Directory domain.
- Password

The password Turbonomic will use to connect to the WinRM service on the Virtual Machine hosting the application.
- Scope

Turbonomic will search for Windows applications on the VMs found within the set scope. This scope can be set to the entire environment, single or multiple clusters, or particular virtual machines. Select this option and choose the scope for application discovery.

NOTE:
A single scope (target) must not contain more than 500 virtual machines. Turbonomic recommends utilization of multiple WMI targets for environments exceeding this limit.
- Domain Name

The Active Directory domain used by Turbonomic in conjunction with the Username for authentication. Leave blank for local accounts.
- Full Validation

When selected, Turbonomic will require every VM in the selected scope to be a valid target. If Turbonomic is unable to authenticate a VM in the scope, the target will not validate and data will not be collected.

Supported Actions

For each discovered entity within the application supply chain, Turbonomic can execute or recommend certain actions, as outlined below.

Entity Type	Action
Applications	<p>Without discovered Guest OS Processes or Application Servers, Turbonomic doesn't perform actions on applications. Instead, it performs resize actions on the host VMs. If host utilization is high enough on the physical machine running the application VM, Turbonomic may also recommend provisioning a new host.</p> <p>For specific application servers, see the individual application entry.</p>
Virtual Machines	<ul style="list-style-type: none"> • Provision additional resources (VMem, VCPU) • Move Virtual Machine • Move Virtual Machine Storage • Reconfigure Storage • Reconfigure Virtual Machine

Monitored Resources

Turbonomic monitors the following resources for the container supply chain:

Entity Type	Commodity
WMI/SNMP Application	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)
Virtual Machine	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) • Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) • Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS • Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

Creating a WMI User Account

The service account Turbonomic uses to connect to WMI should be an Active Directory domain account.

Some enterprises require that the service account does not grant full administrator rights. In that case, you can create a restricted service account for WMI:

1. Add the account to each of the following local groups:

- WinRMRemoteWMIUsers__ (Or Remote Management Users)
- Performance Monitor Users

NOTE:

These groups are standard Windows Server 2012 security groups. If you are using an earlier version of Windows Server and do not see these groups, contact Turbonomic Support for assistance.

2. Grant permissions to the account.

In the WMI Management console, grant the `Enable Account` and `Remote Enable` advanced security permissions to the service account:

- Open the WMI Management console (`wmimgmt`).
- Right-click **WMI Control (Local)** and choose **Properties**.
- Go to the **Security** tab and then click **Security** to display the **Security for Root** dialog.
- Click **Advanced**, select the service account, and click **Edit**.
- Confirm that **This namespace and subnamespace** is selected.
- Select **Enable Account** and **Remote Enable** and click **OK**.

Enabling WMI

Enabling Discovery

To enable Turbonomic discovery of Windows-based Guest Processes, you must configure and enable WMI with WinRM, usually by using a group policy for AD Domains.

When local credentials are provided to Turbonomic, NTLM is the authentication mechanism. When domain credentials are provided, either NTLM or Kerberos can be specified.

NOTE:

If Kerberos is specified, the IP address of the target must be resolvable to a host name using DNS.

While these steps will change slightly between Windows Server versions, the general instructions are:

1. Open an elevated PowerShell prompt on the virtual machine(s) to be discovered as a WMI Target.
2. Enable the WS-Management protocol and set the default configuration for remote management.

Execute `winrm quickconfig`.

3. Enable the WinRM *Negotiate* authentication scheme.

This policy is enabled by default. To enable this policy if it is disabled, execute `winrm set winrm/config/service/auth '@{Negotiate="true"}'`.

4. Set the WinRM `AllowUnencrypted` property to `true` if non-SSL connections are preferred.

This property must be set on both the server and the client. Note that setting this value to `true` does not mean that WMI passes sensitive data in an unencrypted form. It will send only the content of the SOAP messages as plain text:

- Server setting:

```
winrm set winrm/config/service '@{AllowUnencrypted="true"}'
```

- Client setting:

```
winrm set winrm/config/client '@{AllowUnencrypted="true"}'
```

5. Restart the Remote Registry service.

After restarting the Remote Registry service, Turbonomic will discover the WMI targets.

SSL Connections

To connect using SSL, then a certificate must be assigned to WinRM, and an SSL listening socket enabled. Turbonomic will accept self-signed certificates. Complete details can be found in the official WinRM documentation.

Some example SSL-related commands:

- List certificates on the system:

```
ls cert://localmachine/my
```

- Find Thumbprint for the default certificate:

```
7B56C33F029E7605D5C58E5597D0A077FE1D7F1C CN=winsql-server1.corp.mycorp.com
```

- Enable SSL listener:

```
winrm create winrm/config/listener?Address=*+Transport=HTTPS @{Hostname="winsql-server1.corp.mycorp.com";CertificateThumbprint="7B56C33F029E7605D5C58E5597D0A077FE1D7F1C";Port=5985}
```



Database Server Targets

This version of Turbonomic supports the following database servers as targets:

- Oracle
- Microsoft SQL Server
- MySQL

You can add an individual database server as a target, or you can add all matching targets within a given group or scope.

If you set target scope, Turbonomic scans each VM within that group or cluster and tries to connect to a database server over the specified port. Turbonomic adds any database servers it finds as targets.

Supply Chain

Database targets add Database Server entities to the supply chain.

Actions

Turbonomic recommends actions for the database supply chain as follows.

NOTE:

This is a general list of actions for entities discovered for databases. You can see how actions differ per technology in each section that describes adding a specific type of database target.

Entity Type	Action
Database Server	<ul style="list-style-type: none">• Resize Database Memory Recommendation only• Resize Transactions Recommendation only• Resize Connections Recommendation only
Virtual Machines	<ul style="list-style-type: none">• Provision additional resources (VMem, VCPU)• Move Virtual Machine

Entity Type	Action
	<ul style="list-style-type: none"> • Move Virtual Machine Storage • Reconfigure Storage • Reconfigure Virtual Machine

Monitored Resources

Turbonomic monitors the following resources for the database supply chain:

NOTE:

This is a general list of commodities for entities discovered for databases. You can see how actions differ per technology in each section that describes adding a specific type of database target.

Entity Type	Commodity
Database Server	<ul style="list-style-type: none"> • DBMem The memory utilized by the database, as a of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Turbonomic can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM • Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second • Response Time The utilization of the server’s allocated response time Measured in Milliseconds (ms) • Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections • TransactionLog The utilization of the server’s capacity for storage devoted to transaction logs Measured in Kilobytes (KB) • Cache Hit Rate The percentage of accesses that result in cache hits. Measured in a percentage of hits vs total attempts (%)
Virtual Machine	<ul style="list-style-type: none"> • Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) • Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="264 264 1084 394"> • Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) <li data-bbox="264 405 1016 535"> • Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS <li data-bbox="264 546 1047 676"> • Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

SQL Server

Turbonomic supports Microsoft SQL Server 2008 R2, 2012, 2014, and 2016.

NOTE:

SQL Server clusters are not supported by this version of Turbonomic.

Prerequisites

- A user account with SQL permissions including `Connect SQL` and `View Server State` on the database
- The following services must be running, and set to enabled:
 - Net.Tcp Listener Adapter
 - Net.Tcp Port Sharing Service
- TCP/IP is enabled on the port used for Turbonomic discovery
- To enable dynamic port discovery, the port used by the SQL Browser Service

Creating a Service User Account

The user account that Turbonomic uses for its service login must include the following:

- The account must exist in the Security folder within the SQL Server Object Explorer, with the following properties:
 - Enable **SQL Server Authentication**
 - Disable **Enforce password policy**
- The account's security properties must include:
 - Permission to connect to the database through SQL
 - Permission to view the server state

Adding a SQL Server Database to Turbonomic

You can add an individual database server as a target, or you can add all matching targets within a given scope.

To add a database server as a target, you specify:

- **Target Name**
Name displayed in the Turbonomic UI
 - **Username**
Username for the account. This username must not include the AD domain
 - **Password**
Password for the account. This username must not include the AD domain
 - **Scope**
A cluster or group of VMs that host SQL servers
 - **Browsing Service Port**
The port used to communicate with the browsing service. Turbonomic will obtain the SQLServer port for each instance running on each VM in the scope.
 - **SQLServer Port**
The SQL remote port. Turbonomic will use this port if there is no browsing service port specified, or if the browsing service is not available during discovery.
- NOTE:**
Turbonomic will connect to the port specified for the SQL browsing service first. If that connection fails, Turbonomic will connect using the SQLServer Port.
- **Domain Name**
The Active Directory domain used by Turbonomic in conjunction with the Username for authentication. Leave blank for local accounts.
 - **Full Validation**
When selected, Turbonomic will require every VM in the selected scope to be a valid target. If Turbonomic is unable to authenticate a VM in the scope, the target will not validate and data will not be collected.

For additional information about adding SQL Server targets, see the Green Circle article, "[How to Target SQL Servers in VMTurbo \(5.4 and later\)](#)"

Actions

Turbonomic recommends actions for the application supply chain as follows.

Entity Type	Action
Applications	<p>Without discovered Guest OS Processes or Application Servers, Turbonomic doesn't perform actions on applications. Instead, it performs resize actions on the host VMs. If host utilization is high enough on the physical machine running the application VM, Turbonomic may also recommend provisioning a new host.</p> <p>For specific application servers, see the individual application entry.</p>
Virtual Machines	<ul style="list-style-type: none"> • Provision additional resources (VMem, VCPU) • Move Virtual Machine • Move Virtual Machine Storage • Reconfigure Storage • Reconfigure Virtual Machine

Monitored Resources

Turbonomic monitors the following resources for the application server supply chain:

Entity Type	Commodity
Database Server	<ul style="list-style-type: none"> <li data-bbox="261 390 1495 573"> <p>• DBMem</p> <p>The memory utilized by the database, as a of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Turbonomic can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM</p> <li data-bbox="261 583 1333 716"> <p>• Transactions</p> <p>The utilization of the allocated transactions per second for the given virtual application</p> <p>Measured in transactions per second</p> <li data-bbox="261 726 943 858"> <p>• Response Time</p> <p>The utilization of the server’s allocated response time</p> <p>Measured in Milliseconds (ms)</p> <li data-bbox="261 869 1227 1001"> <p>• Connections</p> <p>The utilization of the connection capacity. Only applicable to database servers</p> <p>Measured in Connections</p> <li data-bbox="261 1012 1219 1144"> <p>• TransactionLog</p> <p>The utilization of the server’s capacity for storage devoted to transaction logs</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="261 1155 951 1274"> <p>• Cache Hit Rate</p> <p>The percentage of accesses that result in cache hits.</p> <p>Measured in a percentage of hits vs total attempts (%)</p>
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="261 1299 976 1432"> <p>• Virtual Memory (VMem)</p> <p>The utilization of the VMem allocated to the hosting VM</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="261 1442 959 1575"> <p>• Virtual CPU (VCPU)</p> <p>The utilization of the VCPU allocated to the hosting VM</p> <p>Measured in Megahertz (MHz)</p> <li data-bbox="261 1585 1081 1717"> <p>• Virtual Storage (VStorage)</p> <p>The utilization of the virtual storage capacity allocated for the VM</p> <p>Measured in Kilobytes (KB)</p> <li data-bbox="261 1728 1016 1860"> <p>• Storage Access Operations Per Second (IOPS)</p> <p>The utilization of IOPS allocated for the VStorage on the VM</p> <p>Measured in IOPS</p> <li data-bbox="261 1871 399 1892"> <p>• Latency</p>

Entity Type	Commodity
	The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)



Appendix — Target Configuration

This appendix contains topics that are related to configuring Turbonomic targets.

Cisco Unified Computing System

UCS Blade Provisioning

When managing a UCS fabric target, Turbonomic can provision any blade servers that are installed in a chassis but not currently in operation. If the workload demands more physical compute resources, Turbonomic can automatically direct UCS to provision a blade, or it can recommend that you provision a blade and you can execute the action from the To Do list. To enable this capability, you must perform two basic steps:

- Configure the way UCS and vCenter Server manage information as blades are provisioned

To enable Turbonomic to perform automatic provisioning of UCS blades, you must configure UCS to work with vCenter Server as they both manage resources such as server pools, policies, and Service Profile Templates. This ensures that as Turbonomic directs the UCS Manager to provision a new blade, vCenter Server will recognize that the new physical host is available. Once vCenter Server can recognize the new blade, Turbonomic can direct vCenter Server to move workloads onto the server.

Turbonomic provisions new blades based on the service profiles of operating blades. To enable this, the configuration must include Service Profile Templates, and the operating blades must be bound to these templates.

For information about configuration that enables automated provisioning of blades, see the Cisco Communities post, “UCS PowerTool and VMware PowerCLI automated management of Auto-deploy” at the following location:

https://communities.cisco.com/community/technology/datacenter/ucs_management/cisco_ucs_developed_integrations/blog/2013/09/16/ucs-powertool-and-vmware-powercli-automated-management-of-auto-deploy

This post includes a video that shows “a joint PowerShell integration utilizing both Cisco UCS PowerTool and VMware PowerCLI.” You can also download the scripts from this post and modify them as necessary for your deployment.

- Set the Host Provision action to Automate or Manual for the blade servers. By default, Turbonomic sets the Host Provision action to Recommend.

For any hosts other than blade servers managed by UCS, Turbonomic cannot provision hosts automatically or manually. Instead, it recommends that you provision a host, and you then install the physical machine and wire it to the network.

In a UCS deployment you can have blade servers installed in the chassis and ready to provision. In that case, Turbonomic can direct UCS to provision a new blade to meet workload demands. For these servers, you can set the Host Provision action to Automated or Manual.

NOTE:

It's important that you only set Automated or Manual host provisioning to UCS blades. If you set Host Provision to Automated for other types of hosts, Turbonomic will attempt to perform the action and fail. As a result, you might never see the recommendation to provision a new host of that type.

Turbonomic groups blade servers by chassis. To restrict Automated or Manual settings to blade servers, use this group. You can set the action mode for all blade servers in your environment, or you can set the mode differently for individual chassis.

Enabling Windows Remote Management

Turbonomic communicates with your Hyper-V servers using Web Services Management (WS-Management), which is implemented on Microsoft platforms using Windows Remote Management (WinRM). The following steps show how to enable WinRM on a single host, using the command line.

1. Ensure Windows Firewall is running on the host.

For you to configure WinRM successfully, Windows Firewall must be running on the host. For more information, see the Microsoft Knowledge Base article #2004640 (<http://support.microsoft.com/kb/2004640>).

2. Set up an SPN for the host machine.

The machine must have an SPN of the form, `protocol/host_address`. For example, `WSMAN/10.99.9.2`.

To get a list of SPNs for the machine, execute the following in the command window:

```
setspn -l <vmm-server-name>
```

If there is no valid SPN in the list, create one by running the command:

```
setspn -A protocol/host-address:port where port is optional
```

For example, `setspn -A WSMAN/10.99.9.2:VMM-02`

3. Set up the Windows Remote Management (WinRM) service to run on startup.

Run the `quickconfig` utility to set up the WinRM service. The `quickconfig` utility:

- Configures the WinRM service to auto-start
- Configures basic authentication and disables unencrypted traffic
- Creates a firewall exception for the current user profile
- Configures a listener for HTTP and HTTPS on any IP address
- Enables remote shell access

To run `quickconfig`, log into a command window as Administrator on the host machine. Then execute the following commands:

```
winrm quickconfig
```

Enter `y` to accept the `quickconfig` changes

4. Set permissions on the host machine.

Execute the following commands in the command window to modify the settings made by `quickconfig`:

- To set the memory capacity for remote shells:

```
winrm set winrm/config/winrs @{MaxMemoryPerShellMB="1024" }
```

- To set up an unsecured HTTP connection:

```
winrm set winrm/config/service @{AllowUnencrypted="true" }
```

```
winrm set winrm/config/service/Auth @{Basic="true" }
```

These steps showed you how to enable WinRM for a single host. Some users find the following methods useful for enabling WinRM on multiple hosts:

- [EnablingWinRmViaGlobal Policy Objects \(on page 79\)](#)
- [EnablingWinRMViaPowerShell \(on page 80\)](#)

Enabling WinRM Via Global Policy Objects

You can configure WinRM for all of your Hyper-V targets by creating and linking a Global Policy Object (GPO) within the Hyper-V domain and applying the GPO to all servers.

Follow the steps below to enable Windows Remote Management (WinRM) for your Hyper-V targets.

For additional information, you can also see the Green Circle article, [Application Discovery: GPO for necessary Registry Changes](#)

1. On the AD domain controller, open the Group Policy Management Console (GPMC). If the GPMC is not installed, see <https://technet.microsoft.com/en-us/library/cc725932.aspx>.
2. Create a new Global Policy Object:
 - a. In the GPMC tree, right-click **Group Policy Objects** within the domain containing your Hyper-V servers.
 - b. Choose **Create a GPO in this domain**, and link it here.
 - c. Enter a name for the new GPO and click **OK**.
3. Specify the computers that need access:
 - a. Select the new GPO from the tree.
 - b. On the **Scope** tab, under **Security Filtering**, specify the computer or group of computers you want to grant access. Make sure you include all of your Hyper-V targets.
4. Right-click the new GPO and choose **Edit** to open the Group Policy Management Editor.
5. Configure the WinRM Service:
 - a. In the Group Policy Management Editor, select **Computer Configuration > Policies > Administrative Templates > Windows Components > Windows Remote Management (WinRM) > WinRM Service**.
 - b. Double-click each of following settings and configure as specified:

Allow automatic configuration of listeners (“Allow remote server management through WinRM” on older versions of Windows Server):	Enabled IPv4 filter: *
Allow Basic authentication:	Enabled
Allow unencrypted traffic:	Enabled

6. Configure the WinRM service to run automatically:
 - a. In the Group Policy Management Editor, expand **Computer Configuration > Preferences > Control Panel Settings**.
 - b. Under Control Panel Settings, right-click Services and choose **New > Service**.
 - c. In the New Service Properties window, configure the following settings:

Startup:	Automatic
Service name:	WinRM
Service option:	Service start

7. Enable Windows Remote Shell:
 - a. In the Group Policy Management Editor, select **Computer Configuration > Policies > Administrative Templates > Windows Components > Windows Remote Shell**.
 - b. Double-click the following setting and configure as specified:

Allow Remote Shell Access:	Enabled
----------------------------	---------

8. Add a Windows Firewall exception:
 - a. In the Group Policy Management Editor, open **Computer Configuration > Windows Settings > Security Settings > Windows Firewall > Windows Firewall**.
 - b. Under Windows Firewall, right-click **Inbound Rules** and choose **New > Rule**.
 - c. In the New Inbound Rule Wizard, select **Predefined: Windows Remote Management and Allow the connection**.

The new group policy will be applied during the next policy process update. To apply the new policy immediately, execute the following command at a Powershell prompt:

```
gpupdate /force
```

Enabling WinRM Via PowerShell

Using PsExec, you can run quickconfig on all your Hyper-V servers and change the default settings remotely. PsExec is a component of PsTools, which you can download from <https://technet.microsoft.com/en-us/sysinternals/bb897553.aspx>.

1. Create a text file containing the Hyper-V host names, for example:


```
hp-vx485
hp-vx486
```
2. Since Turbonomic requires changes to the default quickconfig settings, create a batch file containing the following command:

```
@echo off Powershell.exe Set-WSManQuickConfig -Force Powershell.exe Set-Item WSMan:
\localhost\Shell\MaxMemoryPerShellMB 1024
```

NOTE:

If you are connecting via HTTP, you must include the following command:

```
Powershell.exe Set-Item WSMan:\localhost\Service\AllowUnencrypted -Value $True
```

3. Use PsExec to enable WinRM on the remote servers:

```
.\PsExec.exe @<hosts_file_path> -u <username> -p <password> -c <batch_file_path>
```

NOTE:

If you get an error message when executing this command, add the `-h` option (`.\PsExec.exe -h`).

Secure Setup of WSMan

Turbonomic provides a secure option for Hyper-V/VMM Targets which requires that WSMan be set up securely. Use PowerShell to generate a self-signed certificate, and create an HTTPS WinRM listener.

NOTE:

For clustered Hyper-V targets, you do not need to create a listener on each host. Only create a listener on the host that is being added to the "Address" field in the Target Configuration.

To enable secure WSMan on your Hyper-V host:

1. Generate a self-signed certificate using the following command:

```
New-SelfSignedCertificate -CertstoreLocation Cert:\LocalMachine\My -DnsName
"myhost.example.org"
```

2. Find the thumbprint for the certificate for the host:

```
Get-childItem cert:\LocalMachine\My
```

3. Create an HTTPS WinRM listener for the host with the thumbprint you've found:

```
winrm create winrm/config/Listener?Address=*&Transport=HTTPS
'@{Hostname="myhost.example.org"; CertificateThumbprint="THUMBPRINT_YOU_FOUND"}
```

4. Verify the presence of configured listeners:

```
Get-WSManInstance -ResourceURI winrm/config/listener -Enumerate
```

Port Configuration

To support communication between Turbonomic and the API endpoints of your intended target, provide bidirectional access for the following ports:

Target	Port(s)
Vcenter (Monitoring)	80, 443

Target	Port(s)
Vcenter (Tags)	7443, 10443
Hyper-V	5985, 5986
VMM	5985, 5986
VMAX	5988, 5989
Pure	80, 443
HPE 3PAR	5988, 5989, 8080
NetApp	80, 443
Cisco UCS	80, 443

Sample OpenStack SNMP Configuration File for PM Metric Collection

NOTE: This is a sample configuration file for the process found [here \(on page 27\)](#). Please perform those steps before creating this file.

To use this configuration file, you must edit at least the following items:

- <YourLocationName>
- <ContactName>
- <ContactEmailID>

In addition, the line `disk /var/lib/nova/instances 15%` should be edited to `disk / 15%` if the VMs are on local storage.

```
#####
#
# snmpd.conf:
#   An example configuration file for configuring the ucd-snmp snmpd agent.
#
#####
#
# This file is intended to only be as a starting point.  Many more
# configuration directives exist than are mentioned in this file.  For
# full details, see the snmpd.conf(5) manual page.
#
# All lines beginning with a '#' are comments and are intended for you
# to read.  All other lines are configuration commands for the agent.
#####
# Access Control
#####
```

```

# As shipped, the snmpd demon will only respond to queries on the
# system mib group until this file is replaced or modified for
# security purposes. Examples are shown below about how to increase the
# level of access.

# By far, the most common question I get about the agent is "why won't
# it work?", when really it should be "how do I configure the agent to
# allow me to access it?"
#
# By default, the agent responds to the "public" community for read
# only access, if run out of the box without any configuration file in
# place. The following examples show you other ways of configuring
# the agent so that you can change the community names, and give
# yourself write access to the mib tree as well.
#
# For more information, read the FAQ as well as the snmpd.conf(5)
# manual page.

####
# First, map the community name "public" into a "security name"

#      sec.name  source          community
com2sec local    default          public
com2sec mynetwork default        public

####
# Second, map the security name into a group name:

#      groupName      securityModel securityName
group MyRWGroup v1      local
group MyRWGroup v2c     local
group MyRWGroup usm     local
group MyROGroup v1      mynetwork
group MyROGroup v2c     mynetwork
group MyROGroup usm     mynetwork

####
# Third, create a view for us to let the group have rights to:

# Make at least snmpwalk -v 1 localhost -c public system fast again.
#      name          incl/excl      subtree      mask(optional)
view all    included  .1              80

####
# Finally, grant the group read-only access to the systemview view.

access MyROGroup ""      any      noauth      exact  all      none      none
access MyRWGroup ""      any      noauth      exact  all      all       none

#####
# System contact information
#

```

```

# It is also possible to set the sysContact and sysLocation system
# variables through the snmpd.conf file:

#syslocation Unknown (edit /etc/snmp/snmpd.conf)
#syscontact Root <root@localhost> (configure /etc/snmp/snmp.local.conf)

syslocation <YourLocationName>
syscontact <ContactName> <ContactEmailID>

# Example output of snmpwalk:
# % snmpwalk -v 1 localhost -c public system
# system.sysDescr.0 = "SunOS name sun4c"
# system.sysObjectID.0 = OID: enterprises.ucdavis.ucdSnmpAgent.sunos4
# system.sysUpTime.0 = Timeticks: (595637548) 68 days, 22:32:55
# system.sysContact.0 = "Me <me@somewhere.org>"
# system.sysName.0 = "name"
# system.sysLocation.0 = "Right here, right now."
# system.sysServices.0 = 72

#####
# Logging
#

# We do not want annoying "Connection from UDP: " messages in syslog.
# If the following option is commented out, snmpd will print each incoming
# connection, which can be useful for debugging.

dontLogTCPWrappersConnects yes

# -----

#####
# disk checks
#

# The agent can check the amount of available disk space, and make
# sure it is above a set limit.

# disk PATH [MIN=100000]
#
# PATH: mount path to the disk in question.
# MIN: Disks with space below this value will have the Mib's errorFlag set.
#      Default value = 100000.

# Check the / partition and make sure it contains at least 10 megs.
# Note: Use '/' if the VMs are on local storage and
# '/var/lib/nova/instances' if the VMs are on shared storage

#disk / 15%
disk /var/lib/nova/instances 15%

```