



Workload Optimization Manager 3.0.1 Target Configuration Guide

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies are considered un-Controlled copies and the original on-line version should be referred to for latest version.

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Copyright © 2021 Cisco, all rights reserved

Contents

Target Configuration.....	5
Adding and Removing Targets.....	6
Hypervisor Targets.....	8
Microsoft Hyper-V.....	11
Creating A Service User Account.....	14
vCenter Server.....	15
Creating A Service User Account In vCenter.....	20
Other Information Imported from vCenter.....	21
Private Cloud.....	22
OpenStack.....	23
Virtual Machine Manager.....	29
Public Cloud.....	34
Amazon Web Services.....	34
AWS Billing Families.....	39
AWS Billing.....	40
Microsoft Azure.....	42
Microsoft Enterprise Agreement.....	45
Azure Enterprise Agreements.....	47
Storage Targets.....	49
EMC VMAX.....	51
EMC XtremIO.....	53
EMC ScaleIO.....	55
EMC VPLEX.....	57
HPE 3PAR.....	58
NetApp.....	61
Restricted Service Accounts In NetApp.....	63
Pure Storage.....	67
Guest Operating Process Targets.....	70
AppDynamics.....	70
Application Insights.....	74
New Relic.....	76
Dynatrace.....	79
SQL Server.....	82
Application Server Targets.....	85

Apache Tomcat.....	89
JVM Application.....	92
Oracle WebLogic.....	95
IBM WebSphere.....	98
Kubernetes Platform Targets.....	101
Hyperconverged Targets.....	106
Nutanix.....	108
Pinning Nutanix Controller VMs in Generic Hypervisor Mode.....	111
Appendix – Target Configuration.....	113
Cisco Unified Computing System.....	113
Enabling Collection of Memory Statistics: AWS.....	114
Enabling Collection of Memory Statistics: Azure.....	117
Enabling Windows Remote Management.....	117
Enabling WinRM Via Global Policy Objects.....	118
Enabling WinRM Via PowerShell.....	119
Secure Setup of WSMAN.....	119
Port Configuration.....	120
Sample OpenStack SNMP Configuration File for PM Metric Collection.....	121



Target Configuration

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

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

NOTE:

Workload Optimization Manager does not support duplicate instances of the same target. When configuring targets, you must not configure two or more targets to the same address in your environment. For example, you must not configure two different targets to the same AWS account, nor two targets to the same vCenter Server instance.

If you do configure duplicate targets, then actions can fail to execute with an error that begins: `Analysis Exception occurred...`

To resolve this issue, identify the duplicate targets, and delete them until you have only one target for each address.

You can assign instances of the following technologies as Workload Optimization Manager targets:

- Application Servers
 - JVM 6.0+
 - IBM WebSphere Application Server 8.5+
 - Apache Tomcat 7.x, 8.x, and 8.5.x
 - Oracle WebLogic 12c
- Cloud Native
 - OpenShift 3.3+
 - Kubernetes
- Fabric and Network
- Guest OS Processes
 - SNMP
 - WMI: Windows versions 2019, 2016, 2012 / 2012 R2, 2008 R2, 10, 8 / 8.1, and 7
- Hypervisors
 - VMware vCenter 5.1, 5.5, 6.0, 6.5, 6.7, and 7.0+
 - Microsoft Hyper-V 2008 R2, Hyper-V 2012/2012 R2, Hyper-v 2016
- Orchestrator Targets
 - ServiceNow

- Action Script
- Private Cloud Managers
 - Microsoft System Center 2012/2012 R2 Virtual Machine Manager and System Center 2016 Virtual Machine Manager
 - OpenStack Havana – Queens
- Public Cloud Managers
 - Amazon AWS
 - Amazon AWS Billing
 - Microsoft Azure Service Principal
 - Microsoft Enterprise Agreement
- Storage Managers
 - Pure Storage F-series and M-series arrays
 - NetApp Cmode/7mode using ONTAP 8.0+ (excluding AFF and SolidFire)
 - EMC VMAX using SMI-S 8.1+
 - EMC VPLEX Local Architecture with 1:1 mapping of virtual volumes and LUNs
 - EMC ScaleIO 2.x and 3.x
 - EMC XtremIO XMS 4.0+
- Virtual Desktop Infrastructure
 - VMware Horizon

Transport Layer Security Requirements

Workload Optimization Manager 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 Workload Optimization Manager 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 Cisco 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 2048 or greater on your target server. If this does not resolve the issue, please contact Cisco Technical Support.

Adding and Removing Targets

The target services your Workload Optimization Manager 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 Workload Optimization Manager activities you want to perform. For example, the following list shows how vCenter privileges correspond to activities Workload Optimization Manager can perform:

- **Read Only** – Enables Workload Optimization Manager monitoring and simulation (what-if scenarios) only
- **vCenter Administrator** – Enables Workload Optimization Manager monitoring, simulation (what-if scenarios), and automation functions
- **Enable Datastore Browse** – Enabling this property for the account gives Workload Optimization Manager 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.

NOTE:

As you add targets, be sure not to add duplicate entries for the same target.

Typical information you provide includes:

- **Target Type** – Choose among the supported technologies
After you choose the technology, then choose the specific target type for that technology. For example, for Cloud Management, you can choose AWS.
- **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

Removing Targets

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) and/or containers, providing these entities compute and storage resources. When you connect Workload Optimization Manager to hypervisor targets in your environment, Workload Optimization Manager assures application performance by utilizing these resources as efficiently as possible.

Once connected to a hypervisor target, Workload Optimization Manager discovers the VMs, containers, physical machines that host the VMs or containers, datastores that provide storage resources to the physical machines, and virtual datastores that provide storage resources.

As additional targets are added, Workload Optimization Manager will discover the resources belonging to your physical and virtual infrastructure. For example, adding the underlying hardware as part of a UCS and/or storage target will provide additional visibility into the physical infrastructure of your environment. To extend the virtual infrastructure, application server or guest operating process targets can be added.

Workload Optimization Manager represents your environment holistically as a supply chain of resource buyers and sellers, all working together to meet application demand. By empowering buyers (VMs, instances, containers, and services) with a budget to seek the resources that applications need to perform, and sellers to price their available resources (CPU, memory, storage, network) based on utilization in real-time, Workload Optimization Manager maintains your environment within the desired state.

For more information, see "Application Resource Management" in the *User Guide*

Supply Chain

Each hypervisor requires a physical machine (host) and one or more datastores to provide compute and storage resources. Virtual machines (VMs) or containers 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 Workload Optimization Manager extends the supply chain analysis into the components that make up the disk array. For more information, see [Storage Manager Targets \(on page 49\)](#).

Actions

Workload Optimization Manager 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

Entity Type	Action
	<ul style="list-style-type: none"> ■ Reconfigure Storage ■ 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

Workload Optimization Manager 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

Entity Type	Commodity
	<p>Measured in Kilobytes per second (KB/s)</p> <ul style="list-style-type: none"> ■ Swap <p>The utilization of the PM's swap space</p> <p>Measured in Kilobytes (KB)</p> ■ Balloon <p>The utilization of shared memory among VMs running on the host. ESX-only</p> <p>Measured in Kilobytes (KB)</p> ■ 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 Megahertz (MHz)</p>
Storage	<ul style="list-style-type: none"> ■ Storage Amount <p>The utilization of the datastore's capacity</p> <p>Measured in Megabytes (MB)</p> ■ Storage Provisioned <p>The utilization of the datastore's capacity, including overprovisioning.</p> <p>Measured in Megabytes (MB)</p> ■ Storage Access Operations Per Second (IOPS) <p>The summation of the read and write access operations per second on the datastore</p> <p>Measured in Operations per second</p> ■ Latency <p>The utilization of latency on the datastore</p> <p>Measured in Milliseconds (ms)</p>
Datacenter	<p>NOTE: For datacenter entities, Workload Optimization Manager 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) <p>The utilization of the PM's memory reserved or in use</p> <p>Measured in Kilobytes (KB)</p> ■ CPU <p>The utilization of the PM's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> ■ IO <p>The utilization of the PM's IO adapters</p> <p>Measured in Kilobytes per second (KB/s)</p> ■ Net <p>The utilization of data through the PM's network adapters</p> <p>Measured in Kilobytes per second (KB/s)</p> ■ Swap <p>The utilization of the PM's swap space</p> <p>Measured in Kilobytes (KB)</p> ■ Balloon <p>The utilization of shared of memory among VMs running on the host. ESX-only</p> <p>Measured in Kilobytes (KB)</p>

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

Microsoft Hyper-V

If you have a small number of Hyper-V hosts in your environment, you can add them individually as Workload Optimization Manager 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 Workload Optimization Manager 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 Workload Optimization Manager 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 \(on page 29\)](#).

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 \(on page 117\)](#).

For accurate SMB storage calculations, Workload Optimization Manager requires a VMM target.

Prerequisites

- Create a user account that Workload Optimization Manager can use to connect to your Hyper-V servers. See [Creating a Service User Account in Hyper-V \(on page 14\)](#)
- Configure remote management on each Hyper-V server. Refer to [Enabling Windows Remote Management \(on page 117\)](#)
- The time on each Hyper-V host must be in synch with the rest of the managed Hyper-V environment.
- Your Hyper-V environment must not use Server Message Block (SMB) storage.

To manage SMB storage, Workload Optimization Manager 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. Cisco 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**
Workload Optimization Manager 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 \(on page 118\)](#)).
- **Username**
The username of a user account Workload Optimization Manager 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.

NOTE:

If your Hyper-V hosts are running in a Nutanix environment, you must understand pinning a Nutanix Controller VM. For more information, see [Pinning Controller VMs in Generic Hypervisor Mode \(on page 109\)](#).

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.

Workload Optimization Manager 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, Workload Optimization Manager 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

Workload Optimization Manager monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="277 380 1044 485"> Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) <li data-bbox="277 495 930 600"> Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) <li data-bbox="277 611 1044 716"> Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) <li data-bbox="277 726 987 831"> Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS <li data-bbox="277 842 1011 947"> Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> <li data-bbox="277 976 914 1081"> Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) <li data-bbox="277 1092 873 1197"> CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) <li data-bbox="277 1207 776 1312"> IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) <li data-bbox="277 1323 958 1428"> Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) <li data-bbox="277 1438 751 1543"> Swap The utilization of the PM's swap space Measured in Kilobytes (KB)
Storage	<ul style="list-style-type: none"> <li data-bbox="277 1572 779 1677"> Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) <li data-bbox="277 1688 1084 1793"> Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) <li data-bbox="277 1803 1263 1908"> 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

Entity Type	Commodity
	<ul style="list-style-type: none"> ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Datacenter	<p>NOTE: For datacenter entities, Workload Optimization Manager 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 Workload Optimization Manager 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 Workload Optimization 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:

Workload Optimization Manager does not support Restricted User Accounts on Windows 2012 Hyper-V nodes.

To create a restricted service account on your Hyper-V hosts:

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 examples only. If your version of Windows Server does not include these groups, contact Technical Support for assistance.

2. Grant permissions to the service account.

In the WMI Management console, grant the following permissions to the service account:

- Enable Account
- Remote Enable
- Act as Operating System (For Windows 2016)

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 Workload Optimization Manager, you specify a vCenter Server instance as a target. Workload Optimization Manager discovers the associated infrastructure and performs intelligent workload management.

Prerequisites

- A user account Workload Optimization Manager 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. Workload Optimization Manager communicates with each vCenter through a separate API endpoint, but aggregates the data from those vCenters and makes cross-vCenter actions possible.

- **Restricting Workload Optimization Manager Access to Specific Clusters:**

When you add a vCenter target, Workload Optimization Manager 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 Cisco 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 Workload Optimization Manager 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. Cisco 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 Workload Optimization Manager is assigned has the `Profile-driven storage view` permission enabled. This permission is *disabled* in the built-in `readonly` role.

- **Guest Metrics:**

For vCenter Server environments, if you turn on **Enable Guest Metrics** for your target configuration, the metrics that vCenter returns to Workload Optimization Manager can be inaccurate, depending on the version of vCenter Server and VMware Tools that you have running on your VMs.

This target option will instruct that Workload Optimization Manager request from the VC instance that it start collecting and reporting enhanced memory metrics from each virtual machine, using the VMTools installed on each VM. Without this option, Workload Optimization Manager only has access to the "Active" memory for each virtual machine, which is a less accurate representation of the memory that the VM is really using.

For accurate metrics use vCenter Server 6.5U3 or higher with VMware Tools version 10.3.2 or higher on guest VMs. If you run vCenter Server version 7.0 or greater, then it can automatically update your VMware Tools through the user interface. For other versions of vCenter Server, you must manually update your VMware Tools deployments.

For the permissions required for Guest Metrics, see [vCenter Permissions \(on page 17\)](#).

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 Workload Optimization Manager can use to connect to the vCenter Server. Include the domain if required (`<domain>\<username>`).
- **Enable Datastore Browsing**
Enabling datastore browsing enables Workload Optimization Manager to discover wasted storage.
- **Enable Guest Metrics**
Collecting advanced guest memory metrics can increase the accuracy of the VMEM data that Workload Optimization Manager uses for analysis of virtual machines. To enable guest metrics, ensure the following:
 - VMware Tools is installed and running on the target VMs
 - The **Hypervisor VMEM for Resize** vCenter option is active for the VMs discovered by the target.
 - The user account has the `Performance.Modify Intervals` performance privilege
 For more information, see "Hypervisor VMEM for Resize" in the *User Guide* and [vCenter Performance Privileges](#).

vCenter Imported Settings

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

NOTE: Workload Optimization Manager does not import Storage DRS rules at this time.

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

To configure Workload Optimization Manager 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 Workload Optimization Manager 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 Workload Optimization Manager will not initiate VSS → VDS vMotions, vSphere may do so. If this happens, Workload Optimization Manager displays a compliance violation notification.

vCenter Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> Read-only permissions for all entity types
Recommend Actions	<ul style="list-style-type: none"> Read-only permissions for all entity types
Wasted Storage Reporting	<ul style="list-style-type: none"> Datastore > Browse Datastore
Execute VM Move	<ul style="list-style-type: none"> Resource > Migrate Resource > Query Vmotion Resource > Modify Resource Pool Resource > Assign VM to Resource Pool
Execute VM Storage Move	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> Global > Global tags <p>You must also open ports 10443 and 7443 on the target server</p>
Guest Metrics	<ul style="list-style-type: none"> Performance > Modify Intervals

Supported Actions

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

NOTE:

In order to execute cross-VC migrations as a non-admin user, you must have the following permissions enabled for the user account in both origination and destination vCenters:

- Virtual Machine: Edit Inventory, Create From Existing (Move, Register, Remove, Unregister sub-options), Create New
- Datacenter: Reconfigure Datacenter
- Network: Assign Network

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

Workload Optimization Manager monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="342 352 1110 457"> Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) <li data-bbox="342 468 997 573"> Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) <li data-bbox="342 583 1110 688"> Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) <li data-bbox="342 699 1052 804"> Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS <li data-bbox="342 814 1075 919"> Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> <li data-bbox="342 947 980 1052"> Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) <li data-bbox="342 1062 938 1167"> CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) <li data-bbox="342 1178 841 1283"> IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) <li data-bbox="342 1293 1026 1398"> Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) <li data-bbox="342 1409 815 1514"> Swap The utilization of the PM's swap space Measured in Kilobytes (KB) <li data-bbox="342 1524 1227 1629"> Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) <li data-bbox="342 1640 1497 1776"> 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="342 1801 844 1906"> Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB)

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="342 243 1524 352"> <p>■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB)</p> <li data-bbox="342 359 1524 468"> <p>■ 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</p> <li data-bbox="342 474 1524 583"> <p>■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)</p>
Datacenter	<p data-bbox="337 604 1524 659">NOTE: For datacenter entities, Workload Optimization Manager 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="342 686 1524 795"> <p>■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="342 802 1524 911"> <p>■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="342 917 1524 1026"> <p>■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="342 1033 1524 1142"> <p>■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="342 1148 1524 1257"> <p>■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB)</p> <li data-bbox="342 1264 1524 1373"> <p>■ Balloon The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB)</p> <li data-bbox="342 1379 1524 1518"> <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="342 1541 1524 1650"> <p>■ Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="342 1656 1524 1766"> <p>■ CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="342 1772 1524 1837"> <p>■ Storage The utilization of the storage attached to the Provider vDC.</p>

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

Creating A Service User Account In vCenter

The service account you use must have specific permissions on the vCenter. The easiest way to grant Workload Optimization Manager 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 Workload Optimization Manager activities:

Workload Optimization Manager 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 > 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

Workload Optimization Manager Activity	Required Privileges
	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, Workload Optimization Manager:

- 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 them in Workload Optimization Manager.

NOTE:

In vCenter environments, Workload Optimization Manager does not import DRS rules if DRS is disabled on the hypervisor. Further, if Workload Optimization Manager did import an enabled DRS rule and somebody subsequently disables that DRS rule, then Workload Optimization Manager 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, Workload Optimization Manager 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.



Private Cloud

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

Adding a Private Cloud target to Workload Optimization Manager aids in the goal of Application Resource Management (ARM) by preventing bottlenecks within and between your managed entities, resizing virtual machines and containers, prioritizing managed workloads, and reducing storage latency throughout your private cloud, ensuring that applications running in your private cloud always have the resources they need to perform, when they need to perform.

On the private cloud, you can use Workload Optimization Manager 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, Workload Optimization Manager discovers resource partitions that are managed by the cloud manager, the workloads running on these partitions (the VMs and applications), and the supply that hosts workload (the physical machines and storage). Workload Optimization Manager 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

Workload Optimization Manager recommends actions for private cloud infrastructures as follows:

Entity Type	Action
Provider Virtual Datacenters	Workload Optimization Manager 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

Workload Optimization Manager 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 The utilization of the storage attached to the Consumer vDC. Measured in Kilobytes (KB)

OpenStack

To manage your OpenStack® environment, Workload Optimization Manager connects to the Keystone identity service endpoint. Through this connection, Workload Optimization Manager 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.

Workload Optimization Manager has been tested with hypervisors running OpenStack on Red Hat Enterprise Linux and Red Hat RDO.

Prerequisites

Workload Optimization Manager supports OpenStack versions Icehouse through Newton.

You must have the following services installed and enabled, and Workload Optimization Manager 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 Port: 8041

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 Workload Optimization Manager 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 ■ 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

Release	Meters
	<ul style="list-style-type: none"> ■ 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`.

Enabling Reservations

When you add a valid OpenStack target, Workload Optimization Manager can perform its analysis, recommend actions, and perform actions to assure performance and efficiency in your environment. To support these actions, you do not need to perform other configuration.

Workload Optimization Manager also includes reservations – deployment capabilities that can act as workload orchestration, or be integrated into an existing orchestration system (see the Deploy View in the Workload Optimization Manager user interface). These capabilities include:

- Calculate optimal placement for new workload
- Reserve resources for proposed workload, and include the reservations in real-time and planning analysis

An OpenStack user can use the Workload Optimization Manager placement proposals to deploy workload to the optimal locations.

To enable these capabilities, you must install the Cisco Nova Scheduler plugin that matches your version of OpenStack. Cisco delivers the following versions of this plugin on the Cisco Github repository:

- Icehouse
https://raw.githubusercontent.com/turbonomic/nova/stable/icehouse/nova/scheduler/vmt_scheduler.py
- Juno
https://raw.githubusercontent.com/turbonomic/nova/stable/juno/nova/scheduler/vmt_scheduler.py
- Kilo
https://raw.githubusercontent.com/turbonomic/nova/blob/stable/kilo/nova/scheduler/vmt_scheduler.py
- Mitaka / Newton
https://raw.githubusercontent.com/turbonomic/nova/blob/stable/mitaka/nova/scheduler/vmt_scheduler.py

To fetch the scheduler plugin you want, execute the following commands on the Nova controller (substituting the url to the version of scheduler plugin that you want):

```
cd /usr/lib/python2.6/site-packages/nova/scheduler/

curl -O <URL TO CORRECT VERSION OF vmt_scheduler.py>
```

After you execute these commands to add the plugin to your controller, add the following entries to the file, `/etc/nova/nova.conf` under the `[DEFAULT]` section, where you provide the IP address of your Workload Optimization Manager server, and credentials for a Workload Optimization Manager user account that has administrator privileges:

```
scheduler_driver = nova.scheduler.vmt_scheduler.VMTScheduler

vmturbo_rest_uri = <VMTurbo_IPAddress>

vmturbo_username = <VMTurbo_UserName>

vmturbo_password = <VMTurbo_Password>
```

After you restart the Nova scheduler, it can use the plugin to communicate with your Workload Optimization Manager instance.

Collecting Physical Machine Metrics

To collect full metrics from physical machines in your environment and display them in Workload Optimization Manager, 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)\\. (blocks|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 \(on page 121\)](#).
- 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

Workload Optimization Manager 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

Workload Optimization Manager 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

Entity Type	Commodity
	Measured in milliseconds (ms)
Physical Machine	<ul style="list-style-type: none"> <li data-bbox="342 302 1495 411"> <p>■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="342 422 1495 531"> <p>■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="342 541 1495 651"> <p>■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="342 661 1495 770"> <p>■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="342 781 1495 890"> <p>■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB)</p> <li data-bbox="342 900 1495 1010"> <p>■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB)</p> <li data-bbox="342 1020 1495 1142"> <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 Megahertz (MHz)</p>
Datacenter	<p data-bbox="337 1159 1495 1213">NOTE: For datacenter entities, Workload Optimization Manager 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="342 1241 1495 1350"> <p>■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB)</p> <li data-bbox="342 1360 1495 1470"> <p>■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz)</p> <li data-bbox="342 1480 1495 1589"> <p>■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="342 1600 1495 1709"> <p>■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)</p> <li data-bbox="342 1719 1495 1829"> <p>■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB)</p> <li data-bbox="342 1839 1495 1919"> <p>■ Balloon The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB)</p>

Entity Type	Commodity
	<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 Measured in Kilobytes (KB)
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 The utilization of the storage attached to the Consumer vDC. 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. Workload Optimization Manager 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 Workload Optimization Manager 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 \(on page 117\)](#)
- Grant access on every host in the cluster that Workload Optimization Manager will manage to the following WMI namespaces and subspaces:
 - root/virtualization
 - root/virtualization/v2
 - root/microsoft/windows/storage
 - root/microsoft/windows/smb
 - root/wmi
 - root/cimv2
 - root/mscluster

- 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 hot-fix referenced in the Microsoft Knowledge Base article #2842230 (<http://support.microsoft.com/kb/2842230>).
- Apply time synchronization

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

Adding VMM Targets

Workload Optimization Manager uses the address and credentials you provide to discover the VMM target. From the VMM target, Workload Optimization Manager 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.

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

To add VMM targets, select the **Private Cloud > 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 Workload Optimization Manager to communicate over HTTPS. For more information, see [Secure Setup for WSMAN \(on page 119\)](#).

Note that setting a secure connection for VMM does not also set secure connections for the underlying Hyper-V hosts. Any communications between Workload Optimization Manager 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)

Workload Optimization Manager 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 Workload Optimization Manager uses to log into a Hyper-V host must satisfy certain requirements. For more information, see [Creating a Service User Account in Hyper-V \(on page 14\)](#).
- Hosts in the VMM target that you would like to exclude from Workload Optimization Manager.

Workload Optimization Manager will not discover or manage any hosts you enter in this field. These hosts do not participate in the market, and Workload Optimization Manager 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 wild-card 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, Workload Optimization Manager 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, Workload Optimization Manager 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 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. Workload Optimization Manager 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

Workload Optimization Manager 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

Workload Optimization Manager 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)

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) 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, Workload Optimization Manager 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

Entity Type	Commodity
	<p>Measured in Kilobytes per second (KB/s)</p> <ul style="list-style-type: none"> ■ Swap <p>The utilization of the PM's swap space</p> <p>Measured in Kilobytes (KB)</p> ■ Balloon <p>The utilization of shared of memory among VMs running on the host. ESX-only</p> <p>Measured in Kilobytes (KB)</p> ■ 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> ■ CPU <p>The utilization of the Datacenter's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> ■ 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> ■ CPU <p>The utilization of the Datacenter's CPU reserved or in use</p> <p>Measured in Megahertz (MHz)</p> ■ Storage <p>The utilization of the storage attached to the Consumer vDC.</p> <p>Measured in Kilobytes (KB)</p>



Public Cloud

The public cloud provides compute, storage, and other resources on demand. By adding an AWS Billing Target (AWS) or Microsoft Enterprise Agreement (Azure) to use custom pricing and discover reserved instances, you enable Workload Optimization Manager to use that richer pricing information to calculate workload size and RI coverage for your Azure environment.

You can run all of your infrastructure on a public cloud, or you can set up a hybrid environment where you burst workload to the public cloud as needed. Workload Optimization Manager can analyze the performance of applications running on the public cloud, and provision more instances as demand requires. For a hybrid environment, Workload Optimization Manager can provision copies of your application VMs on the public cloud to satisfy spikes in demand, and as demand falls off it can suspend those VMs if they're no longer needed.

With public cloud targets, you can use Workload Optimization Manager to:

- Scale VMs and Databases
- Change storage tiers
- Purchase VM Reservations
- Locate the most efficient workload placement within the hybrid environment, while assuring performance
- Detect unused storage volumes

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

Supply Chain

For public clouds, Workload Optimization Manager discovers Regions and Zones, and adds them as Datacenter Entities. Regions and zones divide the public cloud into managed subsets. A region is typically associated with the geographic location of the cloud resources, and a zone is some division within the region. One region contains multiple zones.

Amazon Web Services

Amazon Web Services (AWS) provides a reliable and scalable infrastructure platform in the cloud. You gain access to this infrastructure through a subscription account with the appropriate organization API permissions. To specify an AWS target, you provide the credentials for that account and Workload Optimization Manager discovers the resources available to you through that account.

In order to discover RI utilization, you must provide Workload Optimization Manager with access to the S3 bucket that contains the AWS Cost and Usage report. Without this access, Workload Optimization Manager's purchase and scale decisions will be made without consideration of this data.

In order to discover all RIs in an AWS billing family, you must add a billing target via the master account. This master account needs cost explorer API access. Without the master account, Workload Optimization Manager will discover only those RIs that are purchased by accounts which have been added as targets.

Adding AWS Targets

For Workload Optimization Manager to manage an AWS account, you provide the credentials via the access key or IAM role that you use to access that account. For information about getting an Access Key for an AWS account, see the Amazon Web Services documentation.

The AWS Target has different target addition requirements based on connection via key or IAM role.

To add an AWS target *without* an IAM Role, specify the following:

- Custom 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.
- Access Key
Provide the **Access Key** for the account you want to manage.
- Secret Access Key
Provide the **Access Key Secret** for the account you want to manage.
- Proxy Host
The IP of your Proxy Host
- Proxy Port
The port required for the proxy above
- Proxy User
The username required for the proxy above
- Proxy Password
The password required for the proxy above

To add an AWS target *with* an IAM Role, specify the following:

- Custom 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.
- IAM Role ARN
Provide the **ARN** for the IAM role used to access the AWS billing information.
- 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.

Firewall and Proxy Access

To use AWS with a proxy or firewall, it must be configured to allow unrestricted access to the following URLs:

Functionality	AWS Endpoint
Price List	api.pricing.us-east-1.amazonaws.com
AWS Bill	{bucket-name}.s3.{region-containing-report}.amazonaws.com

Functionality	AWS Endpoint
CloudWatch	monitoring.{region-id}.amazonaws.com
CloudWatch Events	events.{region-id}.amazonaws.com
CloudWatch Logs	logs.{region-id}.amazonaws.com
EC2	ec2.{region-id}.amazonaws.com
Elastic Load Balancing	elasticloadbalancing.{region-id}.amazonaws.com
IAM	iam.amazonaws.com
Organizations	organizations.{region-id}.amazonaws.com
Relational Database Service (RDS)	rds.{region-id}.amazonaws.com
Resource Groups	resource-groups.{region-id}.amazonaws.com
Service Catalog	servicecatalog.{region-id}.amazonaws.com
S3	s3.{region-id}.amazonaws.com
Storage Gateway	storagegateway.{region-id}.amazonaws.com

Whitelisting AWS Regions

While Workload Optimization Manager is discovering your AWS environment, if it fails to reach one or more AWS regions, then AWS discovery will fail for that target.

There may be policy decisions that prevent Workload Optimization Manager from reaching all AWS regions. For example, if you operate Workload Optimization Manager behind a firewall, you might not be able to reach all the regions that are available to your AWS account. In that case, you need to specify which regions you want Workload Optimization Manager to discover.

For information about how to specify the regions that you want Workload Optimization Manager to discover, contact your support representative.

Cost and Usage Report

In order for Workload Optimization Manager to display month-to-day spend, you must create a cost and usage report in AWS and store it in an S3 bucket.

Enabling Collection of Memory Statistics

For Workload Optimization Manager to collect memory statistics in AWS, you must set up CloudWatch to enable the collection of these statistics on the VMs in your environment. For more information, see [Enabling Collection of Memory Statistics: AWS \(on page 114\)](#).

IAM Authentication

To connect to an AWS target via IAM, Workload Optimization Manager requires an IAM User for authentication. This user must have the following permissions:

NOTE:

Workload Optimization Manager recommends that you set up IAM access via an IAM group that has the necessary permissions. After you create this group, create a user that is a member of it and specify that user for Workload Optimization Manager to access your AWS environment. For more information on IAM Groups, see <http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>.

Workload Optimization Manager also supports logging in to AWS targets via IAM Roles. To enable using Roles, you must run the Workload Optimization Manager software on an EC2 instance in the AWS cloud, and you must have the Workload Optimization Manager instance run as the IAM Role, and connect to the AWS target accounts with the appropriate IAM Role.

For troubleshooting, contact Workload Optimization Manager Technical Support.

AWS Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ■ EC2:DescribeSpotInstanceRequests ■ EC2:DescribeAvailabilityZones ■ EC2:DescribeInstances ■ EC2:DescribeImages ■ EC2:DescribeVolumes ■ EC2:DescribeVolumeStatus ■ EC2:DescribeVolumesModifications ■ EC2:DescribeAddresses ■ EC2:DescribeRegions ■ EC2:DescribeReservedInstances ■ EC2:DescribeReservedInstancesModifications ■ EC2:DescribeSpotPriceHistory ■ EC2:DescribeAccountAttributes ■ CloudWatch:GetMetricStatistics ■ CloudWatch:ListMetrics ■ CloudWatch:GetMetricData ■ RDS:DescribeDBInstances ■ RDS:ListTagsForResource ■ AutoScaling:DescribeAutoScalingGroups ■ ClassicELB:DescribeLoadBalancers ■ ClassicELB:DescribeInstanceHealth ■ AppELB:DescribeTargetGroups ■ AppELB:DescribeTargetHealth ■ Org:DescribeOrganization ■ Org:ListAccounts ■ STS:AssumeRole ■ STS:AssumeRoleWithWebIdentity ■ IAM:GetUser ■ ServiceCatalog:SearchProducts
Action Execution / Recommendation	<ul style="list-style-type: none"> ■ EC2:DescribeInstances ■ EC2:DescribeVolumes ■ EC2:DescribeVolumesModifications ■ EC2:ModifyInstanceAttribute ■ EC2:StopInstances ■ EC2:StartInstances ■ EC2:ModifyVolume ■ EC2:DescribeInstanceStatus ■ EC2:DescribeReservedInstancesOfferings ■ EC2>DeleteVolume ■ AutoScaling:SuspendProcesses ■ AutoScaling:ResumeProcesses ■ AutoScaling:DescribeLaunchConfigurations ■ AutoScaling>CreateLaunchConfiguration ■ AutoScaling>DeleteLaunchConfiguration ■ AutoScaling:UpdateAutoScalingGroup

Workload Optimization Manager Functionality	Required Permissions
	<ul style="list-style-type: none"> ■ ServiceCatalog:DescribeProduct ■ ServiceCatalog:ProvisionProduct ■ ServiceCatalog:DescribeRecord ■ ServiceCatalog:ListLaunchPaths
Savings Plans	<ul style="list-style-type: none"> ■ savingsplans:DescribeSavingsPlans

Actions

Workload Optimization Manager recommends actions for the cloud target supply chain as follows.

Entity Type	Action
Virtual Machines (Cloud)	<ul style="list-style-type: none"> ■ Scale up to template (based on VMem / VCPU) ■ Scale down to template (based on VMem / VCPU) ■ Move Virtual Machine (intra-cloud) <p>NOTE: This is a destructive move- data / applications are not preserved. This action also requires both a Workload Optimization Manager merge policy and the moved VM to be a Linux VM with template configuration.</p>
Zone	<ul style="list-style-type: none"> ■ Start VM on the Zone ■ Suspend VM on the Zone

Monitored Resources

Workload Optimization Manager monitors the following resources for the cloud target supply chain:

Entity Type	Commodity
Virtual Machine (AWS)	<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) ■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS ■ Net Throughput Rate of message delivery over a port Measured in KB/s ■ I/O Throughput The throughput to the underlying storage for the entity Measured in KB/s ■ Latency The utilization of latency allocated for the VStorage on the VM

Entity Type	Commodity
	Measured in milliseconds (ms)
Database Server	<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) Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) I/O Throughput The throughput to the underlying storage for the entity Measured in KB/s Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

AWS Billing Families

The screenshot shows the AWS Billing Families interface. At the top, there is a search bar and a filter icon. Below, a list of 13 targets is shown. Two targets are visible: 'aa.aws.amazon.com' and 'ab.aws.amazon.com', both marked as 'VALIDATED: JAN 28, 202...'. A star symbol next to each indicates a master account. An arrow points to the star with the text 'A star symbol indicates a master account.' Another arrow points to the expand icon with the text 'Expand to see details.'

The 'ab.aws.amazon.com' target is expanded to show details. It lists the account ID 'ABC (010101010101)' as the 'Master account'. Under 'RELATED ACCOUNTS', there are three entries: 'Prod (111111000000)', 'Test (000000111111)', and 'TestABC (121212121212)'. The 'TestABC' entry is greyed out. An arrow points to the greyed-out entry with the text 'A greyed name indicates a member account that you have not configured as a target.' Below the related accounts, the target status is 'Validated' and the last validated time is 'Jan 28, 2021 5:06:41 PM'.

As you configure AWS targets, Workload Optimization Manager discovers AWS accounts that are consolidated into *billing families*. A billing family has one *master* account, and zero or more *member* accounts. By recognizing billing families, Workload

Optimization Manager more accurately calculates cloud investments and savings, and makes more accurate recommendations for RI coverage.

In the Targets user interface, master accounts appear in bold, with a star next to them. You can expand the account entry to see the related member accounts. If you expand the entry for a member account, then the related accounts includes the family master, indicated by a star.

For RI purchases, different accounts in a billing family can share the same RI resources. At the same time, accounts in other billing families cannot use those RIs. This adds flexibility to your RI coverage, while maintaining order over the billing.

In Workload Optimization Manager, if you enable Billing Family Recognition, then you can see the billing family master and member accounts in the Targets user interface, and Workload Optimization Manager can recommend proper RI purchases within the correct billing families.

To enable Billing Family Recognition, ensure the following as you configure your AWS targets:

- Use the proper role for each AWS target

To properly discover billing family information for a target, you must give Workload Optimization Manager credentials for an AWS role that includes the permission, `organizations:DescribeOrganization`. With that permission, Workload Optimization Manager can:

- Discover master accounts and member accounts in different billing families
- Display the account names in the user interface
- Discover billing information for each family and account
- Recommend RI actions that respect billing family boundaries

- Configure targets for the complete billing family

One billing family can consolidate a number of AWS accounts. For Workload Optimization Manager to include these accounts in its analysis, you must configure each one as a separate target. If you do not configure all the accounts in a billing family, then Workload Optimization Manager cannot discover complete billing information for that family, and its analysis will be based on incomplete information.

Workload Optimization Manager displays member accounts that have been configured as targets in regular text. For members that Workload Optimization Manager discovers but have not been configured as targets, Workload Optimization Manager displays their names in grayed text.

If you have enabled Billing Family Recognition, you should keep the following points in mind:

- Billing families can grow

Workload Optimization Manager regularly checks the membership of your billing families. If it discovers a new member account, it adds that account to the list of members. If you have already configured the account as a target, then Workload Optimization Manager includes the new member in its analysis of billing families. If the new member is not already a target, then Workload Optimization Manager lists the new member in grayed text.

- You can configure discounts per billing family

Workload Optimization Manager includes a feature to set a discount for a billing group, and to override that discount for specific template families within that scope. For more information, see "Cloud Discounts" in the *User Guide* and "Discount Override: AWS" in the *User Guide*.

- You might see master accounts that have no member accounts

AWS treats every account you create as a part of a billing family. Assume you created an account, but you had no reason to consolidate its billing with any other accounts. In that case, the account appears in the Workload Optimization Manager user interface as a master account, but it has no member accounts.

AWS Billing

The Workload Optimization Manager AWS Billing target allows users to grant access to a bill which is used to discover billing family relationships. It does not provide access to any operational concern of an AWS account.

Adding an AWS Billing Target

The AWS Billing Target has different target addition requirements based on connection via key or IAM role.

To add an AWS Billing target *without* an IAM Role, specify the following:

- Custom 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.
- Access Key
Provide the **Access Key** for the account you want to manage.
- Secret Access Key
Provide the **Access Key Secret** for the account you want to manage.
- Cost and Usage Report Bucket
Name of the S3 bucket that contains the AWS Cost and Usage report.
- Cost and Usage Report Path
Path in the S3 bucket to the AWS Cost and Usage report.
- Cost and Usage Report Region
Region of the S3 bucket that contains the AWS Cost and Usage report.
- 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.

To add an AWS Billing target *with* an IAM Role, specify the following:

- Custom 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.
- IAM Role ARN
Provide the **ARN** for the IAM role used to access the AWS billing information.
- Cost and Usage Report Bucket
Name of the S3 bucket that contains the AWS Cost and Usage report.
- Cost and Usage Report Path
Path in the S3 bucket to the AWS Cost and Usage report.
- Cost and Usage Report Region
Region of the S3 bucket that contains the AWS Cost and Usage report.
- 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.

AWS Billing Target Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ■ CostExplorer:GetReservationUtilization ■ CostExplorer:GetSavingsPlansUtilization ■ CostExplorer:GetSavingsPlansUtilizationDetails ■ IAM:GetUser ■ Org:DescribeOrganization ■ Org:ListAccounts ■ S3:GetBucketAcl ■ S3:GetObject ■ STS:AssumeRole ■ STS:AssumeRoleWithWebIdentity

Actions

Workload Optimization Manager does not recommend actions for AWS Billing targets. However, the billing information will be used in conjunction with the AWS target to make informed decisions.

Monitored Resources

Workload Optimization Manager does not monitor resources for AWS Billing targets. However, the billing information will be used in conjunction with the AWS target to make informed decisions.

Microsoft Azure

Microsoft Azure is Microsoft's infrastructure platform for the public cloud. You gain access to this infrastructure through a service principal target. To specify an Azure target, you provide the credentials for that account and Workload Optimization Manager discovers the resources available to you through that account.

Azure service principal targets will automatically discover the subscriptions to which the service principal has been granted access in the Azure portal. This in turn will create a derived target for each subscription that inherits the authorization provided by the service principal (e.g. contributor). Derived subscription targets are not directly modifiable but otherwise behave like any other target that may be validated and the inventory discovered.

Prerequisites

- External access via App Registration
- Subscription must be registered with the `Microsoft.Capacity` resource provider. This can be done in the `settings` section of the Azure portal.
- Azure Resource Manager

NOTE:

Workload Optimization Manager will not discover Azure Classic virtual machines, as they do not utilize the Azure Resource Manager.

You must allow at least 30 minutes of discovery time for Workload Optimization Manager to allow full discovery of the Resource Group information.

NOTE:

When you first configure an Azure target, under some circumstances the target has `No Quotas Available`, and so Workload Optimization Manager cannot discover the available templates. This can happen when you initially set up the Azure account and you have not enabled any providers. If this occurs, you can install a single VM in your cloud subscription to make quotas available. Or you can navigate to the Azure Subscriptions Blade and select the subscription you want. Then for the resource providers, register the `Microsoft.Compute` option. For more information, see the following Microsoft article: [Resolve errors for resource provider registration](#).

Azure Required Permissions

- **Action Execution Access**
For action execution access, the user must have the `Owner` or `Contributor` role.
- **Read-Only Access**
For read-only access, ensure the user has the `Subscription - Reader` role and `Storage Account - List Keys` role, **OR** the `Subscription - Reader and Storage - Data Access` role on the storage account where memory metrics are stored.

Creating Client Secret Key and Permission

To create the secret key:

1. Navigate to the `Certificates and Secrets` section of your registered app
2. Click on `+ New client secret`
3. Create a new client secret. The description can be any valid name, and the expiration should be set to `never`

NOTE: Make sure to copy the secret. It **will not** be displayed again

External Access via App registration

The administrator of an Azure Active Directory (Tenant) can register an application with the tenant – This app registration gives an external application access to the tenant's resources. Workload Optimization Manager connects to an Azure target via an App registration.

To create an App registration in your tenant:

- Log into the Azure Management Portal
- Add an App registration to an available tenant – The tenant ID will correspond with the tenant ID that you set for the target.
- From the newly-created App registration blade in the Management Portal, go to the Overview blade
This blade displays the generated Directory (tenant) ID and Application (client) ID for this app.
- From the newly created App registration blade in the Management Portal, go to the Certificates and Secrets blade.
This blade displays previously created certificates and client secrets, as well as provides the ability to create them.

NOTE:

The administrator who creates the App registration must also create a Client secret key. This secret key must be recorded on creation. The administrator can return to the Management Portal to see the Application (client) ID, but the portal only shows the Client secret key once during creation.

For more information, refer to Microsoft's article, [How to: Use the portal to create an Azure AD application and service principal that can access resources](#).

Accessing the Subscriptions

1. Navigate to your Azure subscriptions list.
2. Select the first subscription you want Workload Optimization Manager to manage.
3. Navigate to the Access control for that subscription.
4. Add a role (see permissions prerequisite).
5. Add the application you registered with the tenant.

NOTE:

Repeat steps for each subscription that Workload Optimization Manager will manage.

Locked Storages and Resource Groups

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

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

Accessing Reservations

To manage the use of Azure Reservations, the App registration for this target must have permissions to manage the reservations. In most cases, `Reader` permissions are sufficient. In the case of reservations that are scoped to specific subscriptions, you must provide the app permissions to the reservation order, per the Microsoft article, [Manage Reservations for Azure resources](#). Specifically, review the section titled, *Add or change users who can manage a reservation*.

Adding Azure Targets

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

- **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.
- **Tenant ID**
The ID of the tenant that contains subscriptions to be managed with Workload Optimization Manager.
- **Client App ID**
The Client/App ID of the App Registration that gives Workload Optimization Manager access to resources in your Azure subscription.
- **Client Secret Key**
The secret key for the App Registration.

Firewall and Proxy Access

To use Azure with a proxy or firewall, it must be configured to allow unrestricted access to the following URLs:

- `ratecard.azure-api.net`
- `management.azure.com`
- `login.microsoftonline.com`
- `*.core.windows.net`

Memory Statistics

For Workload Optimization Manager to collect memory statistics in Azure, you must enable the collection of these statistics on the VMs in your environment. You can do this as you deploy your VMs, or you can enable the counters on VMs you have already deployed. For more information, see [Enabling Collection of Memory Statistics: Azure \(on page 117\)](#).

Actions

Workload Optimization Manager recommends actions for the cloud target supply chain as follows.

Entity Type	Action
Virtual Machines (Cloud)	<ul style="list-style-type: none"> ■ Scale up to template (based on VMem / VCPU) ■ Scale down to template (based on VMem / VCPU) ■ Move Virtual Machine (intra-cloud) <p>NOTE: This is a destructive move- data / applications are not preserved. This action also requires both a Workload Optimization Manager merge policy and the moved VM to be a Linux VM with template configuration.</p>

Monitored Resources

Workload Optimization Manager monitors the following resources for the cloud target supply chain:

Entity Type	Commodity
Virtual Machine (Azure)	<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) ■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS ■ I/O Throughput The throughput to the underlying storage for the entity Measured in KB/s ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Database Server	<ul style="list-style-type: none"> ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)
Database	<ul style="list-style-type: none"> ■ Database Transaction Unit (DTU) A DTU represents a mixture of the following performance metrics as a single performance unit for an Azure SQL Database: CPU. Memory. Data I/O and Log I/O.

Microsoft Enterprise Agreement

You can configure Workload Optimization Manager to manage Azure subscriptions within the context of an Enterprise Agreement (EA). An EA target enables Workload Optimization Manager to use custom pricing and discover reserved instances.

When you configure an EA target, Workload Optimization Manager uses that richer pricing information to calculate workload size and RI coverage for your Azure environment.

To enable Workload Optimization Manager management of Azure EA environments, you must configure both an EA target and at least one service principle target. For more information about service principal targets, see [Adding Azure Targets \(on page 42\)](#).

Prerequisites

- Microsoft Azure EA access key
- You must ensure that **DA View Charges** and **AO View Charges** are both enabled in the EA Portal (located under **Manage**). If you just enabled these settings, it can take up to 24 hours for the changes to take effect. For more information, see [Troubleshoot enterprise cost views](#) in the Microsoft Azure documentation.

NOTE: If you perform self-service exchanges for your RIs, Workload Optimization Manager does not discover the new charges for the exchanged RIs through the Azure EA target. To track the charges after you have exchanged RIs, ensure you have an Azure subscription target for the affected scope of Azure workloads, and that subscription has read access to reserved instances information.

- Your firewall must allow access to the `consumption.azure.com:443` address.

Adding Microsoft Enterprise Agreement Targets

To add a Microsoft Enterprise Agreement target, select the **Cloud Management > Microsoft Enterprise Agreement** option on the Target Configuration page and provide the following information:

- Target Name
A user-friendly name that will identify the target
- Enrollment Number
Enterprise Agreement enrollment number (found in your EA admin account at `ea.azure.com`)
- API Key
The API Access Key for the Enterprise Agreement (found in your EA admin account at `ea.azure.com`)
- Proxy Host
The IP address of the proxy server used, if any
- Proxy Port
The port number of the proxy server
- Proxy User
The username of the proxy user used to authenticate
- Proxy Password
The password of the proxy user used to authenticate

When you add the target and it validates, Workload Optimization Manager:

- Recognizes any existing Azure targets in your environment that are part of the EA
- Updates these targets with custom prices from the EA
- Discovers RIs in these targets

Note that this can take up to 24 hours, as target stitching occurs after the next bill processing cycle.

NOTE: Workload Optimization Manager does not generate actions on the EA target specifically, but for the underlying service principal targets. For information about actions and monitored resources for Azure targets, see [Adding Azure Targets \(on page 42\)](#).

Azure Enterprise Agreements

The screenshot displays two target cards in the 'Targets View'. The first card, '13 Targets', is expanded to show an 'Azure-EA' target. This target is linked to an 'EA ENROLLMENT#' and lists 'RELATED SUBSCRIPTIONS' including 'EA - Prod' and 'EA - Test'. The second card, 'core.windows.net', is a 'Service Principal' target that discovers underlying subscriptions. It lists 'RELATED TARGETS' including 'EA: Azure-EA (11111111)', 'EA - Prod', 'EA - Test', and a 'Standalone' target 'EATest'. Annotations on the left explain the roles of these targets.

1) Enterprise Agreement (EA) target

2) EA - Prod is one of the subscriptions in this EA.

3) The Service Principal target (core.windows.net in this example) discovers the underlying subscriptions.

4) Some subscriptions (such as EA - Prod) participate in the EA.

5) Other subscriptions (such as EATest) are standalone or pay-as-you-go.

You can configure Workload Optimization Manager to manage Azure subscriptions within the context of an Enterprise Agreement (EA). An EA defines specific pricing, including the pricing for Reserved Instances (RIs). When you configure an EA target, and set the EA key to your Azure targets, Workload Optimization Manager uses that richer pricing information to calculate workload placement and RI coverage for your Azure environment.

To enable Workload Optimization Manager management of Azure EA environments, you must configure:

- One Microsoft Enterprise Agreement target
- At least one Service Principal target that can discover the underlying Azure subscriptions

For information about Azure targets, see [Microsoft Azure \(on page 42\)](#).

In the Targets View, you can identify the targets related to Azure EA as follows:

- EA Targets

The target that discovers the EA to track pricing and RI information. You can have one EA target per Workload Optimization Manager deployment.
- Azure Subscription Targets

The targets that manage the workloads in your Azure environment. These are discovered by Service Principal targets. Note that not all subscription targets *necessarily* participate in the EA. Expand these entries to see the related Service Principal target. For members of the EA, you can see the related EA target as well.

Subscriptions that do not participate in the EA appear as Standalone targets.

NOTE:

In rare circumstances, you can have a subscription that is not in use – The subscription has no workloads associated with it. In this case, Workload Optimization Manager identifies the subscription as Standalone. This is because the target cannot discover any cost or usage information that would relate the subscription to its EA.

- Service Principal Targets

The Azure target that you configure to discover Azure subscription targets. Expand the entry to see the discovered targets. If you have configured an EA target, the entry lists that as well, along with the EA enrollment number.

Reserved Instances and Azure EA

For Azure environments, Workload Optimization Manager can only discover and use RIs if you have configured a Microsoft Enterprise Account target, and if one or more subscriptions participate in that EA.

To discover and manage RIs in Azure environments, Workload Optimization Manager uses both the EA target and the associated subscription targets. On its own, a subscription target exposes costs for pay-as-you-go pricing. The EA target discovers pricing for the available RI instance types. Workload Optimization Manager combines this information to track:

- RI utilization
- RI coverage
- Virtual machine costs (accounting for RIs)

NOTE:

This release of Workload Optimization Manager does not support RI discovery and management for Classic VMs and Classic Cloud Services. Also, it does not support RI discovery and management for Suppressed Core virtual machines.

Cost Calculations for Azure Environments

To understand the reported costs in your Azure environment, consider these points:

- For targets that participate in the EA, Workload Optimization Manager uses the terms of the given EA, and bases costs on the Offer ID that is effective for the given subscription.
- For VMs in Azure, RI pricing does not include the cost of the OS license. However pricing for on-demand VMs does include the license cost.

NOTE:

For Microsoft Azure EA environments, the projected cost for RI Purchase actions might not match associated costs you find in the Microsoft Pricing Calculator.

Workload Optimization Manager actions can recommend RI purchases. For these recommendations, the action assumes a free Linux OS, so the cost estimate does not include the OS cost. However, The Microsoft Pricing Calculator does include costs for OS licenses. As a result, when you compare the Workload Optimization Manager cost estimates to the values in the Pricing Calculator, it's likely that the two estimates will not match. This difference also affects the Break Even Point that appears in the Recommended RI Purchases chart. Because the recommended purchases do not include Azure costs for OS licenses, the listed Break Even Point can be optimistic.

- For workloads you migrated from on-prem to the Azure cloud, Workload Optimization Manager recognizes Azure Hybrid Benefit (AHUB) savings for RIs and on-demand workloads. The costs you see in Workload Optimization Manager charts include this benefit. However, remember that recommended actions do not include any license cost, so the actions will not reflect any proposed AHUB savings (see above).



Storage Targets

Adding a storage Target enables Workload Optimization Manager to connect to your storage subsystem through a native or SMI-S provider API. Workload Optimization Manager uses the target's API to access and collect information from each of the underlying disk arrays. The information is used to set disk performance characteristics according to the type and capacity of storage, leading to improved workload placement.

Similarly, Workload Optimization Manager determines the relationships between storage controllers and disk arrays, and the location of datastores within those arrays. This information also helps optimize workload placement at a more granular level.

For on-premises applications, this optimization will enable Workload Optimization Manager to make more informed decisions about which storage devices the workloads hosting your applications run on, and assist in assuring application SLO. In the cloud, storage data is handled as part of the public cloud target.

Both virtual machines and containers benefit from this level of optimization. In the case of short-lived containers, Workload Optimization Manager will suggest the best datastore to hold persistent data, and paired with a container or hypervisor target, will select the optimal match of compute and storage resources. For longer-lived containers and virtual machines, each workload will be continually assessed for SLA/SLO, and recommendations to move or resize storages will ensure the continued efficiency of your environment.

The section below describes the storage supply chain. For information on how to add specific storage targets, the resources Workload Optimization Manager 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

Workload Optimization Manager 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

Workload Optimization Manager recommends actions for storage targets as follows.

NOTE:

This is a general list of actions for storage managed by storage controllers. Specific actions Workload Optimization Manager 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, Workload Optimization Manager 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

Workload Optimization Manager 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 of the same type 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)

Entity Type	Commodity
	<p>The summation of the read and write access operations per second on the disk array</p> <p>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.</p> <p>Measured in milliseconds (ms)</p>
Storage Controller	<p>NOTE: Not all targets of the same type 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</p> <p>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</p> <p>Measured in Megabytes (MB)</p>

EMC VMAX

Workload Optimization Manager supports management of VMAX2 and 3 Series storage arrays. The VMAX series is a family of enterprise storage arrays designed for SAN environments. Workload Optimization Manager 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, Workload Optimization Manager discovers all the added disk arrays.

NOTE:

Workload Optimization Manager does not utilize Unisphere. Data is collected exclusively from the SMI-S provider.

Workload Optimization Manager will create Storage Groups based on the SLO levels defined in VMAX3 Targets. By default, Storage vMotion actions will respect these SLO levels based on the configured response time.

Prerequisites

- EMC SMI-S Provider V8.x
- A service account that Workload Optimization Manager 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 the provider address begins with `https`, you must follow the IP with the port used to connect.
- Use Secure Connection

If checked, port 5989 will be used to connect. If unchecked, port 5988 will be used.
- Username

The Username for the SMI-S provider.
- Password

The Password for the SMI-S provider.

Entity Comparison

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

EMC VMAX Name	Workload Optimization Manager 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, Workload Optimization Manager can execute or recommend certain actions, as outlined below.

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

Monitored Resources

When calculating available storage, Workload Optimization Manager 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.

Workload Optimization Manager 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 of the same type 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.

Entity Type	Commodity
	<p>Measured in Megabytes (MB)</p> <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 logical pool. <p>Measured in Operations per second</p> <ul style="list-style-type: none"> ■ Latency The utilization of latency on the logical pool. <p>Measured in milliseconds (ms)</p>
Disk Array	<p>NOTE: Not all targets of the same type 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	<ul style="list-style-type: none"> ■ Storage Amount The utilization of the storage controller's capacity. Measured in Megabytes (MB)

EMC XtremIO

EMC® XtremIO® is a flash-based (SSD) storage solution, designed to push data to applications at higher speeds. The system building blocks are SAN appliances called X-Bricks. A deployment is organized into clusters of X-Bricks, and the clusters are managed by the XtremIO Management Server (XMS).

Workload Optimization Manager connects to X-Bricks through the XMS. The XMS presents a unified view of each connected X-Brick cluster, rather than exposing the individual X-Bricks within each cluster. Within Workload Optimization Manager, each X-Brick cluster displays as a single storage controller with an associated disk array.

The relationship between Storage entities and individual X-Bricks within the cluster is not exposed through the XMS – Workload Optimization Manager cannot make recommendations to move datastores from one X-Brick to another. Additionally, the X-Brick has a fixed form factor – Workload Optimization Manager does not recommend resize actions for disk array or storage controller resources.

Workload Optimization Manager recognizes XtremIO arrays as flash storage and sets the IOPS capacity on discovered arrays accordingly.

Prerequisites

- A service user account on XMS 4.0 or higher – typically the default `xmsadmin` account
Workload Optimization Manager uses this account to connect to the XMS and execute commands through the XtremIO API.

Adding XtremIO Targets

For EMC XtremIO targets, select the **Storage > EMC XtremIO** option on the Target Configuration page and provide the following information:

- Address
The name or IP address of the XtremIO Management Server (XMS).
- Username/Password
Credentials for a user account on the XMS.

After validating the new target, Workload Optimization Manager discovers the connected storage entities. This table compares terms used in XtremIO to those used in Workload Optimization Manager:

XtremIO Name	Workload Optimization Manager Entity
Volume	Storage
XtremIO Cluster	Disk Array
XtremIO Cluster	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 \(on page 49\)](#).

Supported Actions

For each discovered entity, Workload Optimization Manager can execute or recommend certain actions, as outlined below.

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

Monitored Resources

When calculating available storage, Workload Optimization Manager excludes disks devoted to the VNX operating system.

Workload Optimization Manager 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 of the same type 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 of the same type 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 ScaleIO

EMC ScaleIO is an example of Software-Defined Storage for the datacenter. It creates a Virtual SAN overlaying commodity infrastructure that consists of multiple LAN-connected Servers with locally attached commodity Storage. It presents a standard Block Storage interface to Applications accessing the Virtual SAN.

Workload Optimization Manager communicates with the EMC ScaleIO system via the REST API Gateway.

Prerequisites

- EMC ScaleIO 2.x or 3.x
- A service account that Workload Optimization Manager can use to connect to the ScaleIO Gateway.

Adding EMC ScaleIO Targets

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

- Address
The IP or host name of the Gateway.
- Username
The Username for the Gateway service account.

- Password

The Password for the Gateway service account.

Entity Comparison

After validating the new target, Workload Optimization Manager discovers the connected storage entities. This table compares terms used in EMC ScaleIO to those used in Workload Optimization Manager:

EMC ScaleIO Name	Workload Optimization Manager Entity
Volume	Storage
Storage Pool	Disk Array
Protection Domain	Storage Controller

Supported Actions

For each discovered entity, Workload Optimization Manager can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations only
Storage	Provision (Cloning)	Resize (Disabled by default)
Disk Array		Resize Disk Array
Protection Domain		Provision (Cloning)

Monitored Resources

Workload Optimization Manager monitors the following storage resources:

Entity Type	Commodity
Storage	<p>NOTE: Not all targets of the same type 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 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
Disk Array	<p>NOTE: Not all targets of the same type 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)

Entity Type	Commodity
	<p>The summation of the read and write access operations per second on the disk array</p> <p>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.</p> <p>Measured in milliseconds (ms)</p>
Storage Controller	<ul style="list-style-type: none"> ■ Storage Amount <p>The utilization of the storage controller's capacity.</p> <p>Measured in Megabytes (MB)</p>

EMC VPLEX

Workload Optimization Manager supports management of EMC VPLEX virtual storage systems in a local configuration, via the VPLEX API. Currently, Workload Optimization Manager does not support Metro or Geo configurations.

VPLEX is used to aggregate and refine data collected between connected Storage and Hypervisor targets. VPLEX supports one-to-one, one-to-many, and many-to-one relationships between virtual volumes and LUNs. Only one-to-one mapping between virtual volume and LUNs is supported by Workload Optimization Manager.

Prerequisites

- VPLEX Management Server
- Hypervisor target supported by Workload Optimization Manager
- Storage target supported by Workload Optimization Manager

NOTE:

In order for Workload Optimization Manager to make use of the information provided by VPLEX, you must also add the hypervisor and storage layered under it as targets.

Adding EMC VPLEX Targets

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

- Address:
 - The IP or Hostname of the VPLEX Management Server
- Username:
 - The Username for the VPLEX Management Server
- Password:
 - The Password for the VPLEX Management Server
- Port Number:
 - The port number for the remote management connection. The default port number for the VPLEX Management server is 443
- Secure Connection:
 - Select this option to use a secure connection (HTTPS)

NOTE:

The default port (443) uses a secure connection.

Supported Actions

For this target, actions are generated and executed via the underlying storage targets. Workload Optimization Manager will use the enhanced visibility provided by VPLEX to make more intelligent storage decisions- for example, recommending storage vMotion between pools.

HPE 3PAR

HPE 3PAR StoreServ systems use controller nodes to manage pools of storage resources and present a single storage system to consumers. Workload Optimization Manager 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 Workload Optimization Manager can use to connect to the SMI-S and WSPAI providers.

NOTE:

For discovery and monitoring, the Workload Optimization Manager service account must have the `Browse` permission on all monitored domains. To exclude domains from monitoring, the Workload Optimization Manager service account must have no permissions on those domains. For action execution, Workload Optimization Manager 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:

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 SMI-S provider.
In the shell session, execute the command, `showcim`.
3. 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, Workload Optimization Manager discovers the connected storage entities. This table compares terms used in HPE 3PAR to those used in Workload Optimization Manager:

HPE 3PAR Name	Workload Optimization Manager 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 \(on page 49\)](#).

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, Workload Optimization Manager:

- **Discovers each Common Provisioning Group (CPG) as a disk array**
In the Workload Optimization Manager user interface, these disk arrays do not host storage – They appear empty. Workload Optimization Manager 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. Workload Optimization Manager considers this single logical pool when it performs analysis – It can recommend moving storage into or out of the AO group. Also, Workload Optimization Manager 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 Workload Optimization Manager user interface. The display name for this logical pool is the name of the AO Configuration.

Supported Actions

For each discovered entity, Workload Optimization Manager can execute or recommend certain actions, as outlined below.

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

Entity Type	Can Be Automated	Recommendations only
Storage Controller		Provision

Monitored Resources

Workload Optimization Manager 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 of the same type 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 of the same type 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

Entity Type	Commodity
	<ul style="list-style-type: none"> ■ Latency The utilization of latency on the logical pool. Measured in milliseconds (ms)
Storage Controller	<p>NOTE: Not all targets of the same type 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)

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. Workload Optimization Manager connects to these storage controllers to support NetApp targets in 7-Mode and Cluster-Mode (C-Mode).

Prerequisites

- Transport Layer Security (TLS) is enabled
- A service account Workload Optimization Manager can use to connect to the NetApp target

Enabling TLS

Starting with version 5.4, by default Workload Optimization Manager 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 Workload Optimization Manager 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 Cisco 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 2048 or greater on your target server. If this does not resolve the issue, please contact Cisco 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, Workload Optimization Manager 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 Cisco -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 Cisco -application ontapi -authmethod password
```

If you prefer not to grant full administrator rights, see [Creating Restricted Service Accounts In NetApp \(on page 63\)](#)

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 Workload Optimization Manager to use.

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

NetApp Name	Workload Optimization Manager 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 \(on page 49\)](#).

Supported Actions

For each discovered entity, Workload Optimization Manager 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, Workload Optimization Manager 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

Workload Optimization Manager monitors the following storage resources:

NOTE:

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

Entity Type	Commodity
Storage	<ul style="list-style-type: none"> ■ Storage Amount The utilization of the datastore's capacity

Entity Type	Commodity
	<p>Measured in Megabytes (MB)</p> <ul style="list-style-type: none"> ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. <p>Measured in Megabytes (MB)</p> <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 datastore <p>Measured in Operations per second</p> <ul style="list-style-type: none"> ■ Latency The utilization of latency on the datastore <p>Measured in Milliseconds (ms)</p>
Disk Array	<p>NOTE: Not all targets of the same type 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 of the same type 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)

Restricted Service Accounts In NetApp

While Workload Optimization Manager 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 RoleUser -r Workload Optimization ManagerRole
```

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 `RoleUser` with the role and user you created:

```
security login create -role Role-Name -username RoleUser -application ontapi
-authmethod password
```

```
security login create -role Role-Name -username RoleUser -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 Role-Name <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-qtrees-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 Group-Name -r Workload Optimization ManagerRole
```

4. Create a user that is a member of the group.

For example:

```
useradmin user add User-Name -g Group-Name
```

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
 - fcp-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 Role-Name -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 Role-Name -username User-Name -application ontapi -authmethod password
```

Pure Storage

Workload Optimization Manager supports management of the following Pure Storage technologies:

- FlashArray//C
- FlashArray//X

The following technologies are not supported:

- FlashBlade

Because of the improved performance of Pure Storage arrays, Workload Optimization Manager intelligently moves more demanding workloads to Flash-based datastores. Workload Optimization Manager analysis is also able to incorporate Pure Storage de-duplication and compression when recommending actions.

Prerequisites

- A service account Workload Optimization Manager 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 Workload Optimization Manager 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, Workload Optimization Manager discovers the connected storage entities. This table compares terms used in Pure to those used in Workload Optimization Manager:

Pure Name	Workload Optimization Manager Entity
Volume	Storage
Shelf Array	Disk Array
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 \(on page 49\)](#).

Supported Actions

For each discovered entity, Workload Optimization Manager 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

Workload Optimization Manager 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 of the same type 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.

Entity Type	Commodity
	Measured in milliseconds (ms)
Storage Controller	<p>NOTE: Not all targets of the same type 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)



Guest Operating Process Targets

Workload Optimization Manager supports the following Guest Operating Process targets:

- AppDynamics 4.1+
- AppInsights 4.1+
- DynaTrace 1.1+
- NewRelic 4.1+
- SNMP (Simple Network Management Protocol)
- WMI (Windows Management Instrumentation)

Guest Operating Process targets support domains of particular application servers that are controlled by management servers. For such managed domains you will add the management server as a target, and Workload Optimization Manager will discover the managed application servers.

Supply Chain

Guest Operating Process targets add Business Application, Business Transaction, Service, Application Component, Application Server, and Database Server entities to the supply chain. You can navigate to the associated target page to see how these entities map to the target nomenclature.

AppDynamics

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

The Workload Optimization Manager integration with AppDynamics provides a full-stack view of your environment, from application to physical hardware. With information obtained from AppDynamics, Workload Optimization Manager 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

- 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
- DB Monitoring User
- Server Monitoring

To use a custom role, ensure that the role has the `View Server Visibility` permission for both applications and databases.

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 Workload Optimization Manager instance.

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

NOTE: For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications.

Entity Mapping

After validating the new target, Workload Optimization Manager discovers the connected entities. The following table describes the entity mapping between the target and Workload Optimization Manager:

AppDynamics	Workload Optimization Manager
Business Application	Business Application
Business Transaction	Business Transaction
Tier	Service
Node	Application Component
Database	Database Server
Server	Virtual Machine, Physical Machine

Adding an AppDynamics Target

NOTE:

It is possible to monitor certain applications or database servers with both AppDynamics and Workload Optimization Manager, 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 Workload Optimization Manager application target.

To add an AppDynamics instance as a target, specify:

- Hostname or IP Address
The host name or IP Address of the AppDynamics controller instance.
- Port

the port used to connect to the AppDynamics controller. By default, this is set to ports 80 (HTTP) and 443 (HTTPS).

NOTE: For SaaS-based AppDynamics instances, you must use port 443.

- Username@Tenant

Username and account ID with the necessary role(s). The format must be *Username@Tenant*, and the user must have the "Read Only User" and "DB Monitoring User" permissions. This username can be found on the "License > Account" page in AppDynamics. For oAuth authentication, the username must be a user defined as an API Client.

NOTE:

The username and password cannot contain any of the following special characters:

\ / " [] : | < > + = ; , ? * , ' tab space @

- Password

Password for the account used to connect to the AppDynamics instance. For oAuth, this will be the client secret key.

NOTE:

The username and password cannot contain any of the following special characters:

\ / " [] : | < > + = ; , ? * , ' tab space @

- Secure Connection

When checked, Workload Optimization Manager 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.

For more information about creating API client users, see the [AppDynamics Documentation](#).

Actions

Workload Optimization Manager recommends actions for the AppDynamics supply chain as follows.

Entity Type	Action
Application Component	<ul style="list-style-type: none"> ■ Suspend VM Recommendation only. ■ Provision VM Recommendation only.

Monitored Resources

NOTE:

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

Workload Optimization Manager monitors the following resources for the AppDynamics supply chain:

Entity Type	Commodity
Business Transaction	<ul style="list-style-type: none"> ■ Transactions The utilization of the allocated transactions per second for the given business transaction

Entity Type	Commodity
	<p>Measured in transactions per second</p> <ul style="list-style-type: none"> ■ Response Time <p>The utilization of the server's allocated response time</p> <p>Measured in Milliseconds (ms)</p>
Business Applications	<ul style="list-style-type: none"> ■ Transactions <p>The utilization of the allocated transactions per second for the given virtual application</p> <p>Measured in transactions per second</p> <ul style="list-style-type: none"> ■ Response Time <p>The utilization of the server's allocated response time</p> <p>Measured in Milliseconds (ms)</p>
Service	<ul style="list-style-type: none"> ■ Transactions <p>The utilization of the allocated transactions per second for the given service</p> <p>Measured in transactions per second</p> <ul style="list-style-type: none"> ■ Response Time <p>The utilization of the server's allocated response time</p> <p>Measured in Milliseconds (ms)</p>
Application Component	<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> <p>NOTE: This commodity is collected for Java, .NET, and Node.js applications only.</p> <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> <p>NOTE: This commodity is collected for Java, .NET, and Node.js applications only.</p> <ul style="list-style-type: none"> ■ Transactions <p>The utilization of the allocated transactions per second for the given entity</p> <p>Measured in transactions per second</p> <ul style="list-style-type: none"> ■ Heap <p>The utilization of the application component's heap</p> <p>Measured in Kilobytes (KB)</p> <p>NOTE: This commodity is collected for Java, .NET, and Node.js applications only.</p> <ul style="list-style-type: none"> ■ Response Time <p>The utilization of the server's allocated response time</p> <p>Measured in Milliseconds (ms)</p> <ul style="list-style-type: none"> ■ Connections <p>The utilization of the connection capacity. Only applicable to database servers</p> <p>Measured in Connections</p> <p>NOTE: This commodity is NOT collected for WebSphere applications.</p> <ul style="list-style-type: none"> ■ Collection Time <p>The percentage of server uptime spent garbage collecting. Available when the JVM profiler is enabled.</p> <p>Measured in percentage of uptime (%)</p>

Entity Type	Commodity
	<p>NOTE: This commodity is NOT collected for WebSphere applications.</p> <ul style="list-style-type: none"> ■ Threads The utilization of the server’s thread capacity Measured in number of Threads <p>NOTE: This commodity is NOT collected for WebSphere applications.</p>
Database	<ul style="list-style-type: none"> ■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) <p>NOTE: This commodity requires a machine agent present, and database hardware monitoring to be enabled.</p> <ul style="list-style-type: none"> ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) <p>NOTE: This commodity requires a machine agent present, and database hardware monitoring to be enabled.</p> <ul style="list-style-type: none"> ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second ■ 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, Workload Optimization Manager can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM <p>NOTE: This commodity is collected for MS SQL and Mongo databases only.</p> <ul style="list-style-type: none"> ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections <p>NOTE: This commodity is collected for Mongo databases only.</p> <ul style="list-style-type: none"> ■ Transaction Log The utilization of the server’s capacity for storage devoted to transaction logs Measured in Kilobytes (KB) ■ DB Cache Hit Rate The percentage of accesses that result in cache hits. Measured as a percentage of hits vs total attempts (%)

Application Insights

Workload Optimization Manager supports workload management of the application infrastructure monitored by Application Insights via the Workload Optimization Manager integration, which provides a full-stack view of your environment, from application to hosting server. With information obtained from Application Insights, Workload Optimization Manager 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.

Prerequisites

- A valid Application Insights user account with the same permissions detailed for [Microsoft Azure \(on page 42\)](#) targets.

Entity Mapping

After validating the new target, Workload Optimization Manager discovers the connected entities. The following table describes the mapping of entities between the target and Workload Optimization Manager:

Application Insights	Workload Optimization Manager
Application	Application, Application Server

Adding an Application Insights Target

NOTE:

If an application is monitored by Application Insights or Azure, do not add it as a separate Workload Optimization Manager application target.

To add Application Insights as a target, specify:

- **Unique Target Description**
A user-created name that will appear in the Workload Optimization Manager UI.
- **Tenant Name**
The tenant associated to the Azure subscription associated to Application Insights.
- **Azure Subscription ID**
The ID of the Azure subscription with access to the Azure target associated to Application Insights.
- **Client App ID**
The Client/App ID of the App registration that gives Workload Optimization Manager access to resources in your Azure subscription.
- **Client Secret Key**
The secret key for the App registration.
- **Offer ID**
If applicable, the Azure Offer ID related to the Azure subscription.
- **Enrollment Number**
If an Azure EA account, the enrollment number associated to the Azure subscription.

Stitching Application Insights Applications to the Workload Optimization Manager Environment

Each application monitored by Application Insights can have several associated applications, and each of these applications can be deployed on a different hosting server. Likewise, a hosting server may host multiple groups or partial groups of multiple applications.

In order to accurately stitch metrics from Application Insights, the host name or IP address of the hosting server must be discoverable through the Application Insights instance. For most monitored application instances, this is automatic. If the hosting server is not discoverable, Azure tags on the application can be provided to indicate the IP address or hostname.

You can provide an Azure tag in the following name : value format:

```
CWOM-Host-Name : RoleInstance=hostname;RoleInstance=hostname;RoleInstance=hostname;
```

In the preceding example, RoleInstance is the name of the application instance, and hostname is the hosting server. For example: CWOM-Host-Name : cluster-app-a=120.120.120.10;cluster-app-b=120.120.120.11;cluster-app-c=120.120.120.12;

NOTE: If these tags are provided, they will replace any values discovered through the Application Insights API.

Actions

Workload Optimization Manager recommends actions for the Application Insights supply chain as follows.

Entity Type	Action
Application Component	<ul style="list-style-type: none"> ■ Suspend VM Recommendation only. ■ Provision VM Recommendation only.

Monitored Resources

NOTE:

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

Workload Optimization Manager monitors the following resources for the Application Insights supply chain:

Entity Type	Commodity
Application Component	<ul style="list-style-type: none"> ■ 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 entity Measured in transactions per second ■ Response Time The utilization of the server’s allocated response time Measured in Milliseconds (ms)

New Relic

Workload Optimization Manager supports workload management of the application infrastructure monitored by New Relic via the Workload Optimization Manager integration, which provides a full-stack view of your environment, from application instance to host. With information obtained from New Relic, Workload Optimization Manager 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.

Prerequisites

- A valid New Relic user account that includes both APM and infrastructure monitoring.

NOTE: For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications.

Entity Mapping

After validating the new target, Workload Optimization Manager discovers the connected entities. The following table describes the mapping of entities between the target and Workload Optimization Manager:

New Relic	Workload Optimization Manager
APM: Key Transactions	Business Transaction
APM: Application / Service (New Relic One)	Service
APM: Application Instance	Application Component

New Relic	Workload Optimization Manager
-Infra: Database	Database Server
Infra: Host	Virtual Machine

Supported Applications

Workload Optimization Manager supports the following Application languages and associated commodities:

Application Language	Commodities
Java	Virtual CPU, Virtual Memory, Response Time, Transactions, Heap, Collection Time, Threads
.NET	Virtual CPU, Virtual Memory, Response Time, Transactions
Node.js	Virtual CPU, Virtual Memory, Response Time, Transactions, Heap, Collection Time
Python	Virtual CPU, Virtual Memory, Response Time, Transactions
Python	Virtual CPU, Virtual Memory, Response Time, Transactions
PHP	Virtual CPU, Virtual Memory, Response Time, Transactions
GO	Virtual CPU, Virtual Memory, Response Time, Transactions

Supported Databases

Workload Optimization Manager supports the following Database types and commodities:

NOTE: Database commodities are exposed only if the New Relic account used to connect to Workload Optimization Manager has a New Relic Infrastructure Pro subscription.

Database	Commodities
MySQL	Cache Hit Rate
MS SQL	Cache Hit Rate, Virtual Memory, Transactions
OracleDB	Cache Hit Rate, Transactions, Response Time
MongoDB	Virtual Memory, Connections

Adding a New Relic Target

NOTE:

If an application is monitored by New Relic, do not add it as a separate Workload Optimization Manager application target.

To add New Relic as a target, specify:

- Account ID
The New Relic Account ID.
- REST API Key
The REST API Key *provided by the New Relic platform*. For more information, see [Understand New Relic API Keys](#).
- GraphQL API Key
The GraphQL API Key *provided by the GraphQL service*. This is not identical to the REST API Key above. For more information, see [Generate a new API key in the GraphQL Explorer](#).
- EU Region
If checked, Workload Optimization Manager will use the EU API endpoints.
- Proxy Host (Optional)
The IP of the Proxy Host.

- Proxy Port (Optional)
The port required by the proxy.
- Proxy Username (Optional)
The username required by the proxy.
- Proxy Password (Optional)
The password required by the proxy.

Actions

Workload Optimization Manager recommends actions for the New Relic supply chain as follows.

Entity Type	Action
Application Component	<ul style="list-style-type: none"> ■ Suspend VM Recommendation only. ■ Provision VM Recommendation only.

Monitored Resources

NOTE:

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

Workload Optimization Manager monitors the following resources for the New Relic supply chain:

Entity Type	Commodity
Application Component	<ul style="list-style-type: none"> ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) ■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Transactions The utilization of the allocated transactions per second for the given entity 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) ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ Collection Time The percentage of server uptime spent garbage collecting. Available when the JVM profiler is enabled. Measured in percentage of uptime (%) ■ Threads The utilization of the server’s thread capacity

Entity Type	Commodity
	Measured in number of Threads
Database	<ul style="list-style-type: none"> <li data-bbox="285 302 1524 411"> <p>■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB)</p> <li data-bbox="285 411 1524 527"> <p>■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)</p> <li data-bbox="285 527 1524 642"> <p>■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second</p> <li data-bbox="285 642 1524 758"> <p>■ Heap The utilization of the application server's heap Measured in Kilobytes (KB)</p> <li data-bbox="285 758 1524 919"> <p>■ 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, Workload Optimization Manager 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="285 919 1524 1035"> <p>■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second</p> <li data-bbox="285 1035 1524 1150"> <p>■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections</p> <li data-bbox="285 1150 1524 1266"> <p>■ Collection Time The percentage of server uptime spent garbage collecting. Available when the JVM profiler is enabled. Measured in percentage of uptime (%)</p> <li data-bbox="285 1266 1524 1381"> <p>■ Threads The utilization of the server's thread capacity Measured in number of Threads</p> <li data-bbox="285 1381 1524 1497"> <p>■ Transaction Log The utilization of the server's capacity for storage devoted to transaction logs Measured in Kilobytes (KB)</p> <li data-bbox="285 1497 1524 1612"> <p>■ DB Cache Hit Rate The percentage of accesses that result in cache hits. Measured as a percentage of hits vs total attempts (%)</p>

Dynatrace

Workload Optimization Manager supports discovery of applications that are managed by the Dynatrace platform. Workload Optimization Manager 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 Workload Optimization Manager to discover applications through Dynatrace, the applications must be running on VMs in your environment. Also, these VMs must be managed by Workload Optimization Manager targets such as hypervisors or public cloud targets.

- A valid Dynatrace user account.

The account must provide read-only access to the application data that is gathered by the Dynatrace server. It must enable the following switches for the access scope of the Dynatrace API token:

- Access problem and event feed, metrics, and topology
- Read content log
- Read synthetic monitors, locations, and nodes
- Read configuration
- User sessions
- Anonymize user session data for data privacy reasons
- Read audit logs

NOTE: For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications.

Entity Mapping

After validating the new target, Workload Optimization Manager discovers the connected entities. The following table describes the mapping of entities between the target and Workload Optimization Manager:

Dynatrace	Workload Optimization Manager
Application	Business Application
Service	Service
Process	Application Component, Database Server
-	Container
Host	Virtual Machine

Adding a Dynatrace Target

NOTE:

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

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

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 Workload Optimization Manager 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.

Dynatrace Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ■ DataExport ■ metrics.read

Actions

Workload Optimization Manager 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.

Database Servers are discovered only for MySQL and MSSQL databases.

Workload Optimization Manager monitors the following resources for the Dynatrace supply chain:

Entity Type	Commodity
Business Application	<ul style="list-style-type: none"> ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second
Service	<ul style="list-style-type: none"> ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second
Application Component	<ul style="list-style-type: none"> ■ Collection Time The percentage of server uptime spent garbage collecting. Available when the JVM profiler is enabled. Measured in percentage of uptime (%) NOTE: This commodity is collected for Java applications only. NOTE: This commodity is NOT collected for WebSphere applications. ■ Heap The utilization of the application server's heap

Entity Type	Commodity
	Measured in Kilobytes (KB) NOTE: This commodity is collected for Java applications only.
Container	<ul style="list-style-type: none"> ■ VMem The of memory capacity the container utilizes Measured in Megabytes (MB) ■ VCPU The of CPU capacity the container utilizes Measured in Gigahertz (Ghz)
Virtual Machine	<ul style="list-style-type: none"> ■ VMem The of memory capacity the container utilizes Measured in Megabytes (MB) ■ VCPU The of CPU capacity the container utilizes Measured in Gigahertz (Ghz)

SQL Server

Workload Optimization Manager supports the following versions of this target:

NOTE:

SQL Server clusters are not supported by this version of Workload Optimization Manager.

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 Workload Optimization Manager discovery
- To enable dynamic port discovery, the port used by the SQL Browser Service

Creating a Service User Account

The user account that Workload Optimization Manager 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 Workload Optimization Manager

To add an SQL Server target, you add all matching databases within a given scope.

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

- Target Name

- Name displayed in the Workload Optimization Manager 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 group of VMs that host the application.
 - If you set the target scope, Workload Optimization Manager scans each VM within that group or cluster and tries to connect to the target over the specified port. Workload Optimization Manager adds any instances of the target it finds as entities from which metrics are retrieved.
 - The maximum supported size of the group is 500 VMs, and the recommended size is 250 VMs. Adding more VMs to the group may result in delays in discovery and monitoring. If larger numbers of VMs are to be targeted, the recommendation is to split them across smaller groups in multiple targets.
 - Browsing Service Port
 - The port used to communicate with the browsing service. Workload Optimization Manager will obtain the SQLServer port for each instance running on each VM in the scope.
 - SQLServer Port
 - The SQL remote port. Workload Optimization Manager will use this port if there is no browsing service port specified, or if the browsing service is not available during discovery.
- NOTE:**
Workload Optimization Manager will connect to the port specified for the SQL browsing service first. If that connection fails, Workload Optimization Manager will connect using the SQLServer Port.
- Domain Name
 - The Active Directory domain used by Workload Optimization Manager in conjunction with the Username for authentication. Leave blank for local accounts.
 - Full Validation
 - When selected, Workload Optimization Manager will require all VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a VM in the scope, the target will not validate and data will not be collected.

Actions

Workload Optimization Manager recommends actions for the application supply chain as follows.

Entity Type	Action
Applications	<p>Without discovered Guest OS Processes or Application Servers, Workload Optimization Manager 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, Workload Optimization Manager 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

Workload Optimization Manager monitors the following resources for the application server supply chain:

Entity Type	Commodity
Database Server	<ul style="list-style-type: none"> <li data-bbox="256 380 1497 533"> <p>■ 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, Workload Optimization Manager 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="256 541 1497 646"> <p>■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second</p> <li data-bbox="256 655 1497 760"> <p>■ Response Time The utilization of the server’s allocated response time Measured in Milliseconds (ms)</p> <li data-bbox="256 768 1497 873"> <p>■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections</p> <li data-bbox="256 882 1497 987"> <p>■ TransactionLog The utilization of the server’s capacity for storage devoted to transaction logs Measured in Kilobytes (KB)</p> <li data-bbox="256 995 1497 1100"> <p>■ Cache Hit Rate The percentage of accesses that result in cache hits. Measured in a percentage of hits vs total attempts (%)</p>
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="256 1136 1497 1241"> <p>■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB)</p> <li data-bbox="256 1249 1497 1354"> <p>■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)</p> <li data-bbox="256 1362 1497 1467"> <p>■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB)</p> <li data-bbox="256 1476 1497 1581"> <p>■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS</p> <li data-bbox="256 1589 1497 1694"> <p>■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)</p>



Application Server Targets

With the exception of Apache Tomcat and Generic JVM, application server targets support domains of application servers that are controlled by management servers. For such managed domains you can add the management server as a target, and Workload Optimization Manager will discover the managed application servers. You can also add an individual stand-alone application server as a target, or you can add all matching targets within a given scope.

Application Servers are entities that map to management servers for each target. These applications can have application components, which are software components, application code, or units of processing that consume resources to enable the workload to perform its task.

Services are measurable functions which are part of an internal or user-initiated request- for example, a request to update inventory. A service consumes resources from application components.

A virtual machine entity in the context on an application server is the hosting entity of the application, which may reside either on-premises, or in the cloud.

Supply Chain

Application Server targets add Service, Application Component, Application Server, and Database Server entities to the supply chain. The inventory groups applications by server type – You can navigate to each individual application server to see specific details.

Resource Monitoring

To manage application servers, Workload Optimization Manager monitors the resources that affect application performance, as listed in the following table. Note that because WebLogic performs its own tuning of thread pools, Workload Optimization Manager does not monitor threads for WebLogic application servers.

Monitored Resource	WebSphere	WebLogic	JBoss	Tomcat	JVM
Heap	Auto	Auto	Auto	Auto	Auto
Garbage Collection Time	Auto	Auto	Auto	Auto	Auto
Threads	Auto	Auto	Auto	Auto	
Transactions	Auto	Auto	Auto	Auto	
Response Time	Auto	Auto	Auto	Auto	
VMem (from underlying VM)	Auto	Auto	Auto	Auto	

Monitored Resource	WebSphere	WebLogic	JBoss	Tomcat	JVM
VCPU (from underlying VM)					

In addition to monitoring these resources and recommending associated actions, Workload Optimization Manager:

- Monitors garbage collection in the application memory space, and uses that information when recommending actions for heap resizes.
- Discovers JDBC connection pools and through these pools can discover relationships between an application server and a given database server. As a result, Workload Optimization Manager sees the JDBC connection pool as a consumer of database connections, and can recommend resize actions for that connection pool.

Application Server Scaling Policy

After you specify an application server target, Workload Optimization Manager discovers the resource utilization for that server, and recommends appropriate actions. These actions will follow one of the two following scaling policies:

- Provision
Scaling by Provision enables horizontal scaling, where the environment adjusts to increased demand by provisioning new application servers.
- Resize
Scaling by Resize enables vertical scaling, where the environment adjusts to increased demand by resizing the application server, or the VM that hosts the application server.

These scaling policies are mutually exclusive – the affected application servers can scale by Provisioning or by Resizing, but not both. For example, Workload Optimization Manager can recommend resize actions, but if you have set the Scaling Policy to Provision, Workload Optimization Manager will only recommend Provision and Suspend actions.

You set the Scaling Policy in the Policy view – Workload Optimization Manager policies support scope, so you can select specific groups of application servers to set their Scaling Policies as you like. For more information, see "Application Server Actions" in the *User Guide*.

Actions

Workload Optimization Manager recommends actions for the application supply chain as follows.

NOTE:

This is a general list of actions for entities discovered for applications. See specific actions for each target on that targets page.

Entity Type	Action
Service	Workload Optimization Manager does not recommend actions to perform on the service itself, but it does recommend actions to perform on the application components and hosting VMs. For example, assume a service that manages three SQL databases. If a surge in requests degrades performance across all three databases, then Workload Optimization Manager can start a new application component to run another instance of the database application, and bind it to the service. On the other hand, if SQL requests drop off so that the load balancer only forwards requests to two of the databases, Workload Optimization Manager can suspend the dormant database and unbind it.
Application Component	<ul style="list-style-type: none"> ■ Resize Heap This action can only be executed by Workload Optimization Manager when running in a domain controller. Standalone applications will see only a recommendation. ■ Resize Connection Capacity Recommendation only. ■ Suspend VM This action can only be executed by a VM hosted in a vCenter environment. Applications running on other hypervisors will see only a recommendation.
Application Server	<ul style="list-style-type: none"> ■ Resize Heap

Entity Type	Action
	<p>This action can only be executed by Workload Optimization Manager when running in a domain controller. Standalone applications will see only a recommendation.</p> <ul style="list-style-type: none"> ■ Resize Thread Pool Recommendation only. ■ Resize Connection Capacity Recommendation only. ■ Suspend VM This action can only be executed by a VM hosted in a vCenter environment. Applications running on other hypervisors will see only a recommendation. ■ Provision VM Recommendation only.
Database Server	<ul style="list-style-type: none"> ■ Resize Database Memory Recommendation only ■ Resize Connections Recommendation only
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

Workload Optimization Manager monitors the following resources for the application server supply chain:

NOTE:

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

Entity Type	Commodity
Service	<ul style="list-style-type: none"> ■ Transactions The utilization of the allocated transactions per second for the given service Measured in transactions per second
Application Component	<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

Entity Type	Commodity
	<ul style="list-style-type: none"> ■ 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 ■ Remaining Garbage Collection Capacity The percentage of server uptime spent not performing garbage collection Measured in of uptime (%)
Application Server	<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 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

Entity Type	Commodity
	<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>

Apache Tomcat

Workload Optimization Manager supports connecting to individual Tomcat targets. Workload Optimization Manager connects to the Tomcat process as a remote client via remote JMX access. Target configuration includes the port used by the JMX/RMI registry.

Prerequisites

- A valid JMX user account for the Tomcat server.
If Tomcat security is enabled, this must be a Tomcat JMX user with a `readonly` role.
- Tomcat should run on JDK version 7 or 8
- For VMware environments, VMware Tools must be installed on the VM that hosts the Tomcat server. For Hyper-V environments, Hyper-V Integration Services must be installed.

This ensures that the VM hosting the Tomcat server can get its IP address.

- Remote JMX access is enabled through a port that is opened to the firewall.
- Discovered infrastructure.

Workload Optimization Manager discovers Tomcat servers that are running on VMs or containers. The hosting VM or container must already be in your Workload Optimization Manager inventory.

To set the target for a server running on a VM, you must have first discovered the hosting VM through a Hypervisor target. To set the target for a server running in a container, you must have configured container discovery for Tomcat applications.

- For information about container targets, see [Kubernetes Platform Targets \(on page 101\)](#)
- For information about hypervisor targets, see [Hypervisor Targets \(on page 8\)](#)

Configuring JMX Remote Access

Workload Optimization Manager monitors and controls the Tomcat server via JMX Remote access. You must configure a JMX Remote port.

Note that to work with a firewall you should also set the RMI Server port – If you don't set an RMI port, then JMX sets an arbitrary *ephemeral port*, and you can't guarantee that the port will be open to your firewall.

There are two ways to set JMX Remote port on Linux platforms:

- Ports specified as system properties

You can set the port via the system property, `com.sun.management.jmxremote.port`. For example:

```
com.sun.management.jmxremote.port=8050
```

A common way to set this property is to declare it in the `CATALINA_OPTS` system variable – You can set this in the `setenv.sh` script. For example:

```
CATALINA_OPTS=" $CATALINA_OPTS
```

```
-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=8050"
export CATALINA_OPTS
```

Note that this sets the JMX Remote port, but it does not set the RMI Server port – Tomcat startup will specify an ephemeral port for the RMI server.

- Ports specified in a JMX Remote Lifecycle Listener

This listener component fixes the ports used by the JMX/RMI Server. When you configure the listener, you specify both the JMX Remote port and the RMI Server port. This is the preferred method when working with a firewall. For more information, see the Apache Tomcat documentation.

On Windows, the typical installation is with Tomcat as a service. There are two ways to set the JMX Remote port:

- Via `setenv.bat`

Add the property to the `CATALINA_OPTS` environment variable:

```
set "CATALINA_OPTS=%CATALINA_OPTS% -Dcom.sun.management.jmxremote.port=8050"
```

- Use the Tomcat configuration utility (`tomcat7w` or `tomcat8w`)

Set the port with the following command:

```
-Dcom.sun.management.jmxremote.port=8050"
```

To discover the JMX port that is set to an already running Tomcat, you can look in the following locations:

- For Linux platforms, look in the configuration files – Either:
 - `setenv.sh` – Assuming you configured the port by adding it to the `CATALINA_OPTS` environment variable
 - `$CATALINA_HOME/conf/server.xml` – Assuming you configured a JMX Remote Lifecycle Listener in this file
- For Windows platforms, look in:
 - `setenv.bat` – Assuming you configured the port by adding it to the `CATALINA_OPTS` environment variable
 - The Windows registry – Assuming you installed Tomcat as a Windows service using the Tomcat Configuration utility

Adding a Tomcat Target

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

To add a server as a target, specify:

- Target Name
 - Name displayed in the Workload Optimization Manager UI
- Username
 - Username of an account with the Admin role
- Password
 - Password of an account with the Admin role
- Scope
 - A group of VMs that host the application.

If you set the target scope, Workload Optimization Manager scans each VM within that group or cluster and tries to connect to the target over the specified port. Workload Optimization Manager adds any instances of the target it finds as entities from which metrics are retrieved.

The maximum supported size of the group is 500 VMs, and the recommended size is 250 VMs. Adding more VMs to the group may result in delays in discovery and monitoring. If larger numbers of VMs are to be targeted, the recommendation is to split them across smaller groups in multiple targets.

- Port Number

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="367 243 1299 348"> <p>■ Transactions The utilization of the allocated transactions per second for the given application Measured in transactions per second</p> <li data-bbox="367 359 925 464"> <p>■ Heap The utilization of the application server’s heap Measured in Kilobytes (KB)</p> <li data-bbox="367 474 1006 579"> <p>■ Response Time The utilization of the server’s allocated response time Measured in Milliseconds (ms)</p> <li data-bbox="367 590 909 695"> <p>■ Threads The utilization of the server’s thread capacity Measured in Threads</p> <li data-bbox="367 705 1266 810"> <p>■ Connection The utilization of the connection capacity. Only applicable to database servers Measured in Connections</p> <li data-bbox="367 821 1055 926"> <p>■ Garbage Collection The percentage of server uptime spent garbage collecting Measured in of uptime (%)</p>
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="367 951 1023 1056"> <p>■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB)</p> <li data-bbox="367 1066 1023 1171"> <p>■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)</p> <li data-bbox="367 1182 1136 1287"> <p>■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB)</p> <li data-bbox="367 1297 1071 1402"> <p>■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS</p> <li data-bbox="367 1413 1104 1518"> <p>■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)</p>

JVM Application

Workload Optimization Manager supports connecting to individual JVM Applications as targets. Workload Optimization Manager connects to the JVM process as a remote client via remote JMX access. Target configuration includes the port used by the JMX/RMI registry.

Prerequisites

- A valid JMX user account for the JVM application
If JMX security is enabled this must be a JMX user with a `readonly` role

- The application should run on JDK version 7 or 8
- For VMware environments, VMware Tools must be installed on the VM that hosts the application
 - This ensures that the VM hosting the application can get the application's IP address
- Remote JMX access is enabled through a port that is opened to the firewall
- Discovered infrastructure

Workload Optimization Manager discovers JVM applications that are running on VMs or containers. The hosting VM or container must already be in your Workload Optimization Manager inventory.

To set the target for a server running on a VM, you must have first discovered the hosting VM through a Hypervisor target. To set the target for a server running in a container, you must have configured container discovery for JVM applications.

- For information about container targets, see [Kubernetes Platform Targets \(on page 101\)](#)
- For information about hypervisor targets, see [Hypervisor Targets \(on page 8\)](#)

Configuring JMX Remote Access

Workload Optimization Manager monitors and controls the JVM application via JMX Remote access. You must configure a JMX Remote port.

Note that to work with a firewall you should also set the RMI Server port – If you don't set an RMI port, then JMX sets an arbitrary *ephemeral port*, and you can't guarantee that the port will be open to your firewall.

To set the JMX Remote port, pass in the port at the command line when you start your application. For example, to set the port to 8090, start your application with the following options:

```
-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8090
```

Adding a JVM Target

You can add an individual JVM application as a target, or you can add all matching applications within a given scope.

To add an application as a target, specify:

- **Address:** The name or IP address of the VM hosting the JVM application
- **Scope (optional):**

A group of VMs that host the application.

If you set the target scope, Workload Optimization Manager scans each VM within that group or cluster and tries to connect to the target over the specified port. Workload Optimization Manager adds any instances of the target it finds as entities from which metrics are retrieved.

The maximum supported size of the group is 500 VMs, and the recommended size is 250 VMs. Adding more VMs to the group may result in delays in discovery and monitoring. If larger numbers of VMs are to be targeted, the recommendation is to split them across smaller groups in multiple targets.

- **Port Number:** The JMX Remote port
- **Username/Password:** Credentials for a user account with an Admin role

The credentials you provide must match the credentials you specify for JMX login configuration when you start up the application.

If you disable authentication on the application, then you must still provide arbitrary values for **Username** and **Password**. To disable JMX authentication, use the following flags in the command line as you start the application:

```
-Dcom.sun.management.jmxremote.authenticate=false
```

```
-Dcom.sun.management.jmxremote.ssl=false
```

Multiple JVM Targets On Single VM

Note that you can specify targets with different ports, but that run on the same VM (use the same IP address). You can also specify targets via the same scope, but with different ports – This is another way to assign applications running on the same VM to different ports. To do this:

To do this, add the targets in two separate steps. For example, assume you want to add two JVM application targets, and they both run on the VM at 10.10.123.45. One application is on port 123, and the other application is on port 456. To specify these two targets:

- Specify the first target with the following parameters:
 - Scope/Address: VMs_myCluster.mycorp.com
 - Port number: 123
 - Username: AppUser
 - Password: *****

Then click **ADD**.

- Specify the second target with the following parameters:
 - Scope/Address: VMs_myCluster.mycorp.com
 - Port number: 456
 - Username: OtherAppUser
 - Password: *****

Then click **ADD**.

Actions

Workload Optimization Manager recommends actions for the application supply chain as follows.

Entity Type	Action
Virtual Applications	<p>Workload Optimization Manager does not recommend actions to perform on the virtual application itself, but it does recommend actions to perform on the VMs that host bound applications. For example, assume a virtual application that manages three SQL databases. If a surge in requests degrades performance across all three databases, then Workload Optimization Manager can start a new VM to run another instance of the database application, and bind it to the virtual application. On the other hand, if SQL requests drop off so that the load balancer only forwards requests to two of the databases, Workload Optimization Manager can suspend the dormant database and unbind it from the virtual application.</p> <p>Virtual Applications will only be discovered if there is a load balancer configured.</p>
JVM	<ul style="list-style-type: none"> ■ Resize Heap Recommendation only. ■ Resize Thread Pool Recommendation only. ■ Resize Connection Capacity Recommendation only. ■ Suspend VM This action can only be executed by a VM hosted in a vCenter environment. Applications running on other hypervisors will see only a recommendation. ■ Provision VM Recommendation only.
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

Workload Optimization Manager monitors the following resources for the application server supply chain:

Entity Type	Commodity
Virtual Applications	<ul style="list-style-type: none"> ■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second
JVM Application	<ul style="list-style-type: none"> ■ Heap The utilization of the application server's heap Measured in Kilobytes (KB) ■ Garbage Collection The percentage of server uptime spent garbage collecting. Available when the JVM profiler is enabled. Measured in percentage of uptime (%)
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)

Oracle WebLogic

The typical WebLogic deployment is a managed domain with one Administration Server that provides a single point of entry for administration and management of the domain. The domain can include other WebLogic Servers which are the Managed Servers. You set the WebLogic Administration Server as a Workload Optimization Manager target.

For a standalone WebLogic deployment, the single server acts as its own Administration Server – You can set the standalone server as a Workload Optimization Manager target.

WebLogic deployments can include clusters to distribute workload across multiple WebLogic servers. Workload Optimization Manager recommended actions respect the cluster architecture. For example, if you have enabled horizontal scaling for your WebLogic servers, then Workload Optimization Manager can recommend provisioning new servers for a given cluster.

Prerequisites

- A service user account.

To execute actions the service account must have an Admin role. For read-only monitoring and analysis, you can set the target with a more restricted role, but then you will have to execute all recommended actions manually, through the WebLogic interface.

- WebLogic requires both the local and remote ends of the connection be resolvable by DNS. The Workload Optimization Manager IP address and all WebLogic server IP addresses must be resolvable by your local DNS server.
- Discovered infrastructure

Workload Optimization Manager discovers WebLogic servers that are running on VMs or containers. The hosting VM or container must already be in your Workload Optimization Manager inventory.

To set the target for a server running on a VM, you must have first discovered the hosting VM through a Hypervisor target. To set the target for a server running in a container, you must have configured container discovery for WebLogic applications.

- For information about container targets, see [Kubernetes Platform Targets \(on page 101\)](#)
- For information about hypervisor targets, see [Hypervisor Targets \(on page 8\)](#)

Finding the T3 Listen Port

To configure a WebLogic target, you need to know the port that the server listens on for administrative communications. Launch the WebLogic Administration Console:

- Navigate to Domain Structure and display the domain you're interested in
 - Navigate to **Environment > Servers** and select the Domain Administration Server you're setting as a target
- The console displays configuration information for the server, including the T3 listen port.

Adding a WebLogic Target

You can add one or more WebLogic targets in a given scope, including both standalone servers and domain managers.

To add a this target, specify:

- Target Name
Name displayed in the Workload Optimization Manager UI
- Username
Username of an account with the Admin role
- Password
Password of an account with the Admin role
- Scope
A group of VMs that host the application.

If you set the target scope, Workload Optimization Manager scans each VM within that group or cluster and tries to connect to the target over the specified port. Workload Optimization Manager adds any instances of the target it finds as entities from which metrics are retrieved.

The maximum supported size of the group is 500 VMs, and the recommended size is 250 VMs. Adding more VMs to the group may result in delays in discovery and monitoring. If larger numbers of VMs are to be targeted, the recommendation is to split them across smaller groups in multiple targets.

- Port Number
The WebLogic remote port
- Full Validation
When selected, Workload Optimization Manager will require all VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a VM in the scope, the target will not validate and data will not be collected.

Actions

Workload Optimization Manager recommends actions for the application server supply chain as follows:

Entity Type	Action
Service	Workload Optimization Manager does not recommend actions to perform on the service itself, but it does recommend actions to perform on the application components and hosting VMs. For example, assume a service that manages three SQL databases. If a surge in requests degrades performance across all three databases, then Workload Optimization Manager can start a new application component to run another instance of the database application, and bind it to the service. On the other hand, if SQL requests drop off so that the load balancer only forwards requests to two of the databases, Workload Optimization Manager can suspend the dormant database and unbind it.
Application Component	<ul style="list-style-type: none"> ■ Resize Heap This action can only be executed by Workload Optimization Manager when running in a domain controller. Standalone applications will see only a recommendation. ■ Resize Connection Capacity Recommendation only. ■ Suspend VM This action can only be executed by a VM hosted in a vCenter environment. Applications running on other hypervisors will see only a recommendation.
Virtual Machines	<ul style="list-style-type: none"> ■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Reconfigure Storage

Monitored Resources

Workload Optimization Manager monitors the following resources for the application server supply chain:

Entity Type	Commodity
Service	<ul style="list-style-type: none"> ■ Transactions The utilization of the allocated transactions per second for the given service Measured in transactions per second
Application Component	<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

Entity Type	Commodity
	<p>The utilization of the connection capacity. Only applicable to database servers Measured in Connections</p> <ul style="list-style-type: none"> Remaining Garbage Collection Capacity The percentage of server uptime spent not performing garbage collection 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 Measured in Kilobytes (KB) Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

IBM WebSphere

The typical WebSphere deployment is a cell of WebSphere servers, controlled by a Deployment Manager. A cell makes up a managed domain that incorporates multiple VMS that host managed application servers. The Deployment Manager is a WebSphere instance that provides a single point of entry for the managed domain.

NOTE:

When adding a WebSphere Deployment Manager as a target, you must ensure that the name of each WebSphere node is resolvable to an IP address by the Workload Optimization Manager instance.

You may need to make changes to your DNS or the file `/etc/resolv.conf` on the Workload Optimization Manager instance to make it aware of the domain names in use in your environment.

To configure the WebShpere installation, you can use the WebShpere Integrated Solutions Console. This is a client that exposes configuration settings including the SOAP port and the PMI settings.

To manage the servers in an installation, WebSphere uses the Performance Monitoring Infrastructure (PMI). Each WebSphere server runs a PMI service that collects performance data from the various application server components. Workload Optimization Manager uses PMI for monitoring and control of the WebSphere installation.

Prerequisites

- The PMI service set to monitor at the Basic level or greater
- A service user account

To execute actions the service account must have an Administrator role. For read-only monitoring and analysis, you can set the target with a more restricted role (Monitor), but then you will have to execute all recommended actions manually, through the WebSphere interface.

- Discovered infrastructure

Workload Optimization Manager discovers WebSphere servers that are running on VMs or containers. The hosting VM or container must already be in your Workload Optimization Manager inventory.

To set the target for a server running on a VM, you must have first discovered the hosting VM through a Hypervisor target. To set the target for a server running in a container, you must have configured container discovery for WebSphere applications.

- For information about hypervisor targets, see [Kubernetes Platform Targets \(on page 101\)](#)
- For information about container targets, see [Hypervisor Targets \(on page 8\)](#)

Finding the SOAP Connector Address

To configure a WebSphere target, you need to know the port that the server listens on for administrative communications. Launch the WebSphere Administration Console:

- Navigate to System **Administration > Deployment Manager**
- Under **Additional Properties**, click **Ports**

The entry for `SOAP_CONNECTOR_ADDRESS` gives the currently set port number.

Adding a WebSphere Target

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

To add a server as a target, specify:

- **Target Name**
Name displayed in the Workload Optimization Manager UI
- **Username**
Username of an account with the Admin role
- **Password**
Password of an account with the Admin role
- **Scope**
A group of VMs that host the application.

If you set the target scope, Workload Optimization Manager scans each VM within that group or cluster and tries to connect to the target over the specified port. Workload Optimization Manager adds any instances of the target it finds as entities from which metrics are retrieved.

The maximum supported size of the group is 500 VMs, and the recommended size is 250 VMs. Adding more VMs to the group may result in delays in discovery and monitoring. If larger numbers of VMs are to be targeted, the recommendation is to split them across smaller groups in multiple targets.

- **Port Number**
The WebSphere remote port
- **Full Validation**
When selected, Workload Optimization Manager will require all VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a VM in the scope, the target will not validate and data will not be collected.

Actions

Workload Optimization Manager recommends actions for the application server supply chain as follows:

Entity Type	Action
Service	Workload Optimization Manager does not recommend actions to perform on the service itself, but it does recommend actions to perform on the application components and hosting VMs. For example, assume a service that manages three SQL databases. If a surge in requests degrades performance across all three databases, then Workload Optimization Manager can start a new application component to run another instance of the database application, and bind it to the service. On the other hand, if SQL requests drop off so that the load balancer only forwards requests to two of the databases, Workload Optimization Manager can suspend the dormant database and unbind it.
Application Component	<ul style="list-style-type: none"> ■ Resize Heap This action can only be executed by Workload Optimization Manager when running in a domain controller. Standalone applications will see only a recommendation. ■ Resize Connection Capacity Recommendation only. ■ Suspend VM

Entity Type	Action
	This action can only be executed by a VM hosted in a vCenter environment. Applications running on other hypervisors will see only a recommendation.
Virtual Machines	<ul style="list-style-type: none"> ■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Reconfigure Storage

Monitored Resources

Workload Optimization Manager monitors the following resources for the application server supply chain:

Entity Type	Commodity
Service	<ul style="list-style-type: none"> ■ Transactions The utilization of the allocated transactions per second for the given service Measured in transactions per second
Application Component	<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 ■ Remaining Garbage Collection Capacity The percentage of server uptime spent not performing garbage collection 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 Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)



Kubernetes Platform Targets

Containers support separation of concerns in a way similar to virtual machines, but allow greater flexibility and use far less overhead. Containers may be deployed singly (uncommon) or within a cluster containing multiple nodes. A single container can implement a complete application, or one container can implement a single process that contributes to a larger, distributed application.

With container targets, Workload Optimization Manager can discover entities related to container platforms in your environment. In support of Application Resource Management (ARM), Workload Optimization Manager will recommend resize, placement, and scale actions, and display discovered entities.

Prerequisites

- Running Kubernetes 1.8+ Cluster

Adding a Kubernetes Target

Detailed steps for adding a Kubernetes target can be found on the [KubeTurbo Wiki](#).

Supply Chain

Workload Optimization Manager adds several entities to the supply chain: Services, Containers, Container Pods, Container Specs, Workload Controllers, Namespaces, Volumes, and VirtualMachines. Each entity represents key components of your containerized application running in Kubernetes.

Actions

Workload Optimization Manager recommends actions for the Kubernetes container platform supply chain as follows.

Entity	Action
Service	None No actions are recommended at this level of the supply chain. Instead, actions that affect the service are generated and executed on underlying entities.
Application Component	Suspend Application components are suspended due to a node (virtual machine) suspension APM Actions Application components may also receive other actions as part of APM integration related to those use cases. For example, a <code>Resize Heap</code> action from an underlying AppDynamics integration. See the Target Configuration Guide documentation for the appropriate technology to discover what actions may be available.

Entity	Action		
Container	<p>Resize Container Up/Down With <code>Merged Actions</code> enabled, individual Container actions will be recommend only and the resize will be reflected as an action on the Workload Controller entity.</p> <p>Suspend Containers are suspended due to a node (virtual machine) suspension</p>		
Container Pod	<p>Move Pod Pods will be moved across nodes (Virtual Machines).</p> <p>Suspend Container Pods are suspended due to a node (virtual machine) suspension</p>		
Container Spec	<p>None No actions are recommended at this level of the supply chain. This entity maintains the history of all replicas, or instances of pods for this container specification.</p>		
Workload Controller	<p>Resize Container With <code>Merged Actions</code> enabled, this is a single resize action representing all resize actions for containers associated to a specific workload controller.</p>		
Namespace	<p>None No actions are recommended at this level of the supply chain. Namespace Quotas are constraints to container resizing actions.</p>		
Virtual Machine (Node)	<p>Provision Additional Resources The following resources may be provisioned:</p> <ul style="list-style-type: none"> ■ VMem ■ VCPU ■ VMem Requests ■ VCPU Requests ■ Number of Consumers <p>Suspend Nodes (virtual machines) may be suspended.</p> <p>Infrastructure-dependent Actions Depending on the technology the node is stitched to for underlying infrastructure, there may be additional actions:</p> <table border="1" style="width: 100%;"> <tr> <td style="width: 30%;">On-prem VMware:</td> <td> <ul style="list-style-type: none"> ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine </td> </tr> </table>	On-prem VMware:	<ul style="list-style-type: none"> ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine
On-prem VMware:	<ul style="list-style-type: none"> ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine 		
Volume	<p>None No actions are recommended at this level of the supply chain. These entities will be stitched to public cloud storage volumes.</p>		

Monitored Resources

Workload Optimization Manager monitors the following resources for the Kubernetes container platform supply chain:

Entity	Commodity
Service	Response Time

Entity	Commodity
	<p>Response time of the service, measured in ms. This commodity is populated via APM or DIF integrations.</p> <p>Transactions</p> <p>Transaction utilization, measured in transactions per second. This commodity is populated via APM or DIF integrations.</p>
Application Component	<p>Various Commodities</p> <p>The commodities monitored and the values received for those commodities at the application component level is dependent on the APM integration used. See the Target Configuration Guide documentation for the appropriate technology to discover what data will be reported.</p>
Container	<p>VMem</p> <p>The virtual memory utilized by the container against the memory limit (if no limit is set, then node capacity is used). Measured in Megabytes (MB)</p> <p>VCPU</p> <p>The virtual CPU utilized by the container against the CPU limit (if no limit is set, then node capacity is used). Measured in Megahertz (Mhz)</p> <p>VMemRequest</p> <p>If applicable, the virtual memory utilized by the container against the memory request. Measured in Megabytes (MB)</p> <p>VCPURquest</p> <p>If applicable, the virtual CPU utilized by the container against the CPU request. Measured in Megahertz (Mhz)</p>
Container Pod	<p>VMem</p> <p>The virtual memory utilized by the pod against the node physical capacity. Measured in Megabytes (MB)</p> <p>VCPU</p> <p>The virtual CPU utilized by the pod against the node physical capacity. Measured in Megahertz (Mhz)</p> <p>VMemRequest</p> <p>The virtual memory request allocated by the pod against the node allocatable capacity. Measured in Megabytes (MB)</p> <p>VCPURquest</p> <p>The virtual CPU request allocated by the pod against the node allocatable capacity. Measured in Megahertz (Mhz)</p> <p>VMemRequestQuota</p> <p>If applicable, The amount of virtual memory request the pod has allocated against the namespace quota. Measured in Megabytes (MB)</p> <p>VCPURquestQuota</p> <p>If applicable, The amount of virtual CPU request the pod has allocated against the namespace quota. Measured in Megahertz (Mhz)</p> <p>VMemLimitQuota</p> <p>If applicable, The amount of virtual memory limit the pod has allocated against the namespace quota. Measured in Megabytes (MB)</p> <p>VCPULimitQuota</p> <p>If applicable, The amount of virtual CPU limit the pod has allocated against the namespace quota. Measured in Megahertz (Mhz)</p>

Entity	Commodity
Container Spec	<p>VMem The virtual memory historically utilized by any containers run for this workload against the memory limit (if no limit is set, then node capacity is used). Measured in Megabytes (MB)</p> <p>VCPU The virtual CPU historically utilized by any containers run for this workload against the CPU limit (if no limit is set, then node capacity is used). Measured in Megahertz (Mhz)</p> <p>VMemRequest If applicable, the virtual memory historically utilized by any containers run for this workload against the memory request. Measured in Megabytes (MB)</p> <p>VCPURquest If applicable, the virtual CPU historically utilized by any containers run for this workload against the CPU request. Measured in Megahertz (Mhz)</p>
Workload Controller	<p>VMemRequestQuota If applicable, The amount of virtual memory request the pod has historically allocated for this workload against the namespace quota. Measured in Megabytes (MB)</p> <p>VCPURquestQuota If applicable, The amount of virtual CPU request the pod has historically allocated for this workload against the namespace quota. Measured in Megahertz (Mhz)</p> <p>VMemLimitQuota If applicable, The amount of virtual memory limit the pod has historically allocated for this workload against the namespace quota. Measured in Megabytes (MB)</p> <p>VCPULimitQuota If applicable, The amount of virtual CPU limit the pod has historically allocated for this workload against the namespace quota. Measured in Megahertz (Mhz)</p>
Namespace	<p>VMemRequestQuota The total amount of virtual memory request for all pods allocated to the namespace against the namespace quota. Measured in Megabytes (MB)</p> <p>VCPURquestQuota The total amount of virtual CPU request for all pods allocated to the namespace against the namespace quota. Measured in Megahertz (Mhz)</p> <p>VMemLimitQuota The total amount of virtual memory limit for all pods allocated to the namespace against the namespace quota. Measured in Megabytes (MB)</p> <p>VCPULimitQuota The total amount of virtual CPU limit for all pods allocated to the namespace against the namespace quota. Measured in Megahertz (Mhz)</p>
Virtual Machine (Node)	<p>VMem The virtual memory utilized by the node against the memory allocated to the hosting virtual machine. Measured in Megabytes (MB)</p> <p>VCPU The virtual CPU utilized by the node against the CPU allocated to the hosting virtual machine. Measured in Megahertz (Mhz)</p> <p>VMemRequest The total amount of virtual memory allocated to pods with memory request against the allocatable capacity of the node. Measured in Megabytes (MB)</p>

Entity	Commodity
	<p>VCPURquest</p> <p>The total amount of virtual CPU allocated to pods with CPU request against the allocatable capacity of the node. Measured in Megahertz (Mhz)</p> <p>Number Consumers</p> <p>The total number of pods running on the node against the maximum number of pods allowed. Measured in Pods (#)</p> <p>Infrastructure-dependent Commodities</p> <p>Depending on the technology the node is stitched to for underlying infrastructure, there may be additional commodities, or more granular data reported to existing commodities. See the Target Configuration Guide documentation for the appropriate technology to discover what data will be reported.</p>



Hyperconverged Targets

A hyperconverged target is a service that unites compute, network and storage access into a cohesive system. When you connect Workload Optimization Manager to hyperconverged targets, it will monitor the performance and resource consumption of your hyperconverged infrastructure to maintain application performance while utilizing resources as efficiently as possible.

As part of this process, Workload Optimization Manager will stitch information from the hyperconverged target to the associated hypervisor and fabric targets, supporting Application Resource Management (ARM) and providing deeper insight into the state of the hardware and information related to the entities in the supply chain. Combined with application server targets, this information will support a top-down, application-driven approach to managing your environment.

Monitored Resources

Workload Optimization Manager monitors the following resources for the hyperconverged supply chain, once stitched to your hypervisor and other associated targets:

NOTE: The entities visible in the supply chain depend on what supplemental targets have been added in addition to the hyperconverged target.

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)
Blade	<ul style="list-style-type: none">■ Net The utilization of data through the Blade's network adapters

Entity Type	Commodity
	<p>Measured in Kilobytes per second (KB/s)</p> <ul style="list-style-type: none"> ■ Treated as a Physical Machine of the underlying Hypervisor (see below) CPU, Mem, etc.
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)
I/O Module	<ul style="list-style-type: none"> ■ NetThroughput Rate of message delivery over a port Measured in Megabits per second (Mb/s)
Fabric Interconnect	<ul style="list-style-type: none"> ■ NetThroughput Rate of message delivery over a port Measured in Mb/s ■ PortChannel Amalgamation of ports with a shared net throughput and utilization Measured in Mb/s
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

Entity Type	Commodity
	Measured in Operations per second <ul style="list-style-type: none"> ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	<p>NOTE: Not all targets of the same type 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 <p>The utilization of the Disk Array's capacity.</p> Measured in Megabytes (MB) ■ Storage Provisioned <p>The utilization of the Disk Array's capacity, including overprovisioning.</p> Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) <p>The summation of the read and write access operations per second on the disk array</p> Measured in Operations per second ■ Latency <p>The utilization of latency, computed from the latency of each device in the disk array.</p> Measured in milliseconds (ms)

Nutanix

Nutanix products provide hyperconverged platforms that include VM hosting and a distributed storage fabric. The platform presents storage in two tiers – Local HDD storage and server-attached flash (hot storage).

Nutanix environments may include:

- One or more Nutanix appliances

An appliance contains up to four server nodes.
- Nutanix nodes

Servers that expose compute and storage resources. Each node provides local HDD and hot storage. Nodes combine to form a unified cluster that pools resources.
- Controller VMs

Each node includes a Controller VM that manages the node's resources within the cluster pool. To minimize storage latency, the Controller VM keeps the most frequently accessed data in the hot storage.

Workload Optimization Manager supports management of Nutanix fabrics, where the supply chain treats a Nutanix Storage Pool as a disk array. Workload Optimization Manager recognizes Nutanix storage tiers when calculating placement of VMs and VStorage. In addition, Workload Optimization Manager can recommend actions to scale flash capacity up or down by adding more hosts to the cluster, or more flash drives to the hosts.

To specify a Nutanix target, provide the Cluster External IP address. This is a logical IP address that always connects to one of the active Controller VMs in the cluster. In this way, you can specify a Nutanix target without having to specify an explicit Controller VM.

NOTE:

The Controller VM must remain *pinned* to its host machine – You must not move the Controller VM to a different host. If the Nutanix cluster uses the Nutanix Acropolis OS to manage VMs, Workload Optimization Manager automatically pins the Controller VMs. However, if you use vCenter Server or Hyper-V to manage VMs on the hosts, you must configure a group to pin the Controller VMs. For more information, see [Pinning Nutanix Controller VMs in Generic Hypervisor Mode \(on page 111\)](#).

Prerequisites

- A service account with cluster administrator rights on the Nutanix cluster(s) for action execution. For entity discovery, a minimum of READ access is required.

Finding the Cluster External IP Address

To configure a Nutanix target, provide the Cluster External IP address for the given Nutanix cluster.

The Cluster External IP address is a logical IP that resolves to the cluster's Prism Element Leader. If the Prism Element Leader fails, then the Cluster External IP address will resolve to the newly elected Prism Element Leader.

To find this IP address, open the Web Console (the Prism Element) on the cluster and navigate to the **Cluster Details** view. In this view you can see the **Cluster External IP** address. If there is no IP address specified, you can specify the address at this time. For more information, see the Nutanix documentation.

Operating Modes

A Nutanix node is a server that hosts VMs – In this sense the node functions as a hypervisor. A cluster of nodes can host VMs using the following Hypervisor technologies:

- Nutanix Acropolis
The native Nutanix host platform, which combines software-defined storage with built-in virtualization.
- VMware ESXi
- Microsoft Hyper-V

Workload Optimization Manager supports Nutanix cluster management in the Generic Hypervisor Mode (ESXi or Hyper-V). In this mode you:

- Add each Hyper-V host or vCenter as a hypervisor target – This enables VM workload control for the respective hypervisor technologies
- Specify the Nutanix Cluster External IP address as the target address – This adds the cluster as a Storage Controller target to enable Workload Optimization Manager storage control

Controller VM Pinning

Each Nutanix node hosts a Controller VM that runs the Nutanix software and manages I/O for the hypervisor and all VMs running on the host. Each Controller VM must remain on its host node –The Controller VM must be *pinned* to that host, and must not be moved to any other host.

For more information about how to pin the Controller VM, see [Pinning Nutanix Controller VMs in Generic Hypervisor Mode \(on page 111\)](#).

Adding Nutanix Targets

NOTE:

This describes how to add a Nutanix cluster to Workload Optimization Manager as a target. Before you add the cluster as a target, you should know which operating mode you intend. If you want Standalone mode, then you will have to enable that operating mode after adding the cluster. If you want Generic Hypervisor mode, then you will have to add the hypervisors as targets after you have added the Nutanix cluster as a target. For more information, see [Hypervisor Targets \(on page 8\)](#).

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

- Address
The Cluster External IP address for the Nutanix cluster.
- Port Number
The listening port of the cluster.
- Secure Connection
When Workload Optimization Manager will use a secure connection.
- Username/Password
Credentials for an account on the Nutanix cluster with sufficient privileges.

After validating the new target, Workload Optimization Manager discovers the connected storage entities. This table compares terms used in Nutanix to those used in Workload Optimization Manager:

Nutanix Name	Workload Optimization Manager Entity
Container	Storage
Storage Pool	Disk Array
Nutanix Cluster	Storage Controller

Supported Actions

For each discovered entity, Workload Optimization Manager can execute or recommend certain actions, as outlined below.

Entity Type	Can Be Automated	Recommendations only
VM (a Nutanix VM)	Move (Host), Resize Resize actions require the VM to power down, and power back on again. NOTE: Workload Optimization Manager can automate VMotion to hosts, but for storage moves on Nutanix Workload Optimization Manager only supports the Recommend action mode.	
Datastore ("Storage")	Provision, Resize Up, Resize Down, Suspend	Move
Disk Array		
Storage Controller		Provision

Monitored Resources

Workload Optimization Manager 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	NOTE: Not all targets of the same type 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. <ul style="list-style-type: none"> ■ Storage Amount The utilization of the Disk Array's capacity.

Entity Type	Commodity
	<p>Measured in Megabytes (MB)</p> <ul style="list-style-type: none"> ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. <p>Measured in Megabytes (MB)</p> <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 <p>Measured in Operations per second</p> <ul style="list-style-type: none"> ■ Latency The utilization of latency, computed from the latency of each device in the disk array. <p>Measured in milliseconds (ms)</p>
Storage Controller	<p>NOTE: Not all targets of the same type 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)

Pinning Nutanix Controller VMs in Generic Hypervisor Mode

Each Nutanix node hosts a Controller VM that runs the Nutanix software and manages I/O for the hypervisor and all VMs running on the host. Each Controller VM must remain on its host node –The Controller VM must be *pinned* to that host, and must not be moved to any other host.

For a cluster in Generic Hypervisor mode (using vCenter or Hyper-V hypervisors), you must use Workload Optimization Manager policies to pin the Controller VMs to their respective nodes. To do this, you will create a dynamic group of Nutanix Controller VMs, and then disable move actions for all members of this group.

To pin the Controller VMs:

1. Create a group of Controller VMs.

In Workload Optimization Manager you can create dynamic groups based on VM name – All VMs with matching names automatically belong to the group. Nutanix uses the following naming convention for Control VMs:

NTNX-`<SerialNumber>`-A-CVM, where `<SerialNumber>` is the serial number of the Controller VM.

You can create a dynamic group that automatically includes these Nutanix controller VMs. (For complete instructions on creating groups, see "Creating Groups" in the *User Guide*.)

- Create a new group
In Workload Optimization Manager navigate to **Settings > Groups** and create a new group.
- Set the group type to **Dynamic**
- Add a filter to match VMs by their names
Add a filter that uses the regular expression, `NTNX.*CVM`. This regular expression will match the Nutanix Controller VMs.

Be sure to save the group. All the Nutanix Controller VMs will automatically become members of this group.

2. Disable moves for all VMs in this group.

To do this, create an automation policy for the group and disable actions. (For complete instructions to create these policies, see "Creating Scoped Automation Policies" in the *User Guide*.)

- In Workload Optimization Manager go to the Policy > Action > VM view
- Set the scope to the group you made
In the **Scope** column, expand **My Groups** and select the group you just made.
- Disable moves for this group
In the **Parameter** column under **Action Mode Settings**, set the value to **Disabled**. This will override the global action mode.
- Save the action mode settings
Be sure to click **Apply Settings Change**.



Appendix – Target Configuration

This appendix contains topics that are related to configuring Workload Optimization Manager targets.

Cisco Unified Computing System

UCS Blade Provisioning

When managing a UCS fabric target, Workload Optimization Manager can provision any blade servers that are installed in a chassis but not currently in operation. If the workload demands more physical compute resources, Workload Optimization Manager 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 Workload Optimization Manager 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 Workload Optimization Manager 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, Workload Optimization Manager can direct vCenter Server to move workloads onto the server.

Workload Optimization Manager 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, Workload Optimization Manager sets the Host Provision action to Recommend.

For any hosts other than blade servers managed by UCS, Workload Optimization Manager 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, Workload Optimization Manager can direct UCS to provision a new blade to meet workload demands. For these servers, you can set the Host Provision action to Automatic or Manual.

NOTE:

It's important that you only set Automatic or Manual host provisioning to UCS blades. If you set Host Provision to Automatic for other types of hosts, Workload Optimization Manager 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.

Workload Optimization Manager groups blade servers by chassis. To restrict Automatic 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 Collection of Memory Statistics: AWS

So Workload Optimization Manager can collect memory statistics in your AWS environment, you must set up your VMs to publish statistics via CloudWatch. The steps to do this are different depending on the OS running on your VMs.

AWS Memory Statistics: Linux VMs

To enable memory statistics on Linux VMs, you must set up your VMs to publish metrics to CloudWatch. You must meet the following requirements:

- Your Linux image must have the SSM Agent installed

By default, Linux AMIs dated 2017.09 and later include an installed SSM Agent.

- Access to the CloudWatch service

Your AWS Instance must have internet access or direct access to CloudWatch so it can push data to CloudWatch.

- Access to AWS System Manager

The user account must include an attached AmazonEC2RoleforSSM Policy. As a minimum, the policy must have the following permissions:

- AmazonEC2ReadOnlyAccess
- AmazonS3ReadOnlyAccess
- AmazonRDSReadOnlyAccess
- AmazonEC2RoleforSSM
- CloudWatchFullAccess
- AmazonSSMFullAccess

In addition, you must install the CloudWatch agent on your VMs. To configure the agent, add the following JSON as the EC2 Parameter Store:

```
{
  "metrics": {
    "namespace": "Windows System",
    "append_dimensions": {
      "InstanceId": "${aws:InstanceId}"
    },
    "aggregation_dimensions" : [ ["InstanceId"] ],
    "metrics_collected": {
      "Memory": {
        "measurement": [
          {"name": "% Committed Bytes In Use", "rename": "MemoryUsed", "unit": "Bytes"},
          {"name": "Available Bytes", "rename": "MemoryAvailable", "unit": "Bytes"}
        ],
        "metrics_collection_interval": 60
      },
      "Paging File": {
```

```

    "measurement": [
      {"name": "% Usage", "rename": "paging_used"}
    ],
    "metrics_collection_interval": 60,
    "resources": [
      "*"
    ]
  }
}
}
}

```

To enable additional dimensions (such as InstanceType and ImageId), add the following:

```

{
  "agent": {
    "metrics_collection_interval": 60,
    "logfile": "/opt/aws/amazon-cloudwatch-agent/logs/amazon-cloudwatch-agent.log"
  },
  "metrics": {
    "namespace": "custom_namespace",
    "metrics_collected": {
      "mem": {
        "measurement": [
          {"name": "used", "rename": "MemoryUsed"},
          {"name": "mem_available", "rename": "MemoryAvailable"}
        ]
      }
    },
    "append_dimensions": {
      "InstanceId": "${aws:InstanceId}",
      "InstanceType": "${aws:InstanceType}",
      "ImageID": "${aws:ImageId}"
    }
  }
}

```

AWS Memory Statistics: Windows VMs

To enable memory statistics on Windows VMs, you must enable AWS SSM:

- Ensure the AWS SSM Agent is installed on your VMs
- Create an IAM Role that supports AWS SSM. The role must have the following permissions at a minimum:
 - AmazonEC2ReadOnlyAccess
 - AmazonS3ReadOnlyAccess
 - AmazonRDSReadOnlyAccess
 - AmazonEC2RoleforSSM
 - CloudWatchFullAccess
 - AmazonSSMFullAccess
- Assign the IAM role to the VMs that you want to manage.

After you enable AWS SSM, you must deploy the following JSON file to configure CloudWatch.

```

"IsEnabled": true,

```

```

    "EngineConfiguration": {
      "Components": [
        {
          "FullName": "AWS.EC2.Windows.CloudWatch.PerformanceCounterComponent.PerformanceCounterInputCompon
ent,AWS.EC2.Windows.CloudWatch",
          "Id": "PerformanceCounter",
          "Parameters": {
            "CategoryName": "Memory",
            "CounterName": "Committed Bytes",
            "DimensionName": "InstanceId",
            "DimensionValue": "{instance_id}",
            "InstanceName": "",
            "MetricName": "MemoryUsed",
            "Unit": "bytes"
          }
        },
        {
          "FullName": "AWS.EC2.Windows.CloudWatch.PerformanceCounterComponent.PerformanceCounterInputCompon
ent,AWS.EC2.Windows.CloudWatch",
          "Id": "PerformanceCounter2",
          "Parameters": {
            "CategoryName": "Memory",
            "CounterName": "Available Bytes",
            "DimensionName": "InstanceId",
            "DimensionValue": "{instance_id}",
            "InstanceName": "",
            "MetricName": "MemoryAvailable",
            "Unit": "bytes"
          }
        },
        {
          "FullName": "AWS.EC2.Windows.CloudWatch.CloudWatch.CloudWatchOutputComponent,AWS.EC2.Windows.Clo
udWatch",
          "Id": "CloudWatch",
          "Parameters": {
            "AccessKey": "",
            "NameSpace": "Windows System",
            "Region": "",
            "SecretKey": ""
          }
        }
      ],
      "Flows": {
        "Flows": [
          "(PerformanceCounter,PerformanceCounter2), CloudWatch"
        ]
      },
      "PollInterval": "00:05:00"
    }
  }

```

Enabling Collection of Memory Statistics: Azure

For Workload Optimization Manager to collect memory statistics in Azure, you must enable the collection of these statistics on the VMs in your environment. You can do this as you deploy your VMs, or you can enable the counters after the fact on VMs you have already deployed. For each VM, open the Azure Portal and navigate to Diagnostics Settings. Then enable the metrics for your VMs:

To enable the collection of memory statistics in Azure environments, open the Azure Portal, and then navigate to **Diagnostic Settings**. Then enable the metrics for your VMs:

- For Windows VMs

Navigate to **Performance Counters**, display **Basic**, and enable the performance counters for the VM.

- For Linux VMs

For supported Linux versions, Azure automatically deploys the Linux Diagnostics Extension v2.3 to gather these metrics. Refer to Microsoft Azure documentation for supported Linux OS versions. For unsupported versions, you can enable the statistics manually:

1. Set **Status** to ON.
2. For **Storage Account**, specify the storage that will retain the metric data.
3. Enable **Basic Metrics** and then click **Save**.
4. Navigate to **Metrics** in the Azure Portal and enable the metrics to collect.

Enabling Windows Remote Management

Workload Optimization Manager 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 \(on page 118\)](#)
- [EnablingWinRMViaPowerShell \(on page 119\)](#)

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.

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 Cisco 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

Workload Optimization Manager 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 Workload Optimization Manager and the API endpoints of your intended target, provide bidirectional access for the following ports:

NOTE:

This list may include targets not available to your version or distribution of Workload Optimization Manager.

Target	Port(s)
Vcenter (Monitoring)	80, 443
Vcenter (Tags)	7443, 10443
Vcenter (Kubernetes)	88
Hyper-V	5985, 5986
Hyper-V (Kubernetes)	88
RHEV	80, 443
XenCenter	80, 443
Vcloud Director	80, 443
CloudStack	80, 443
OpenStack	5000, 8774, 8776, 8777
VMM	5985, 5986
CloudFoundry	80, 443
VMAX	5988, 5989
Dell Compellent	5988, 5989
Nutanix	9440
XtremIO	80, 443
Pure	80, 443
HPE 3PAR	5988, 5989, 8080
NetApp	80, 443
NetScaler	80, 443
Cisco UCS	80, 443
sFlow	80, 443
Netflow	80, 443
Arista	80, 443

Target	Port(s)
WebSphere	8880
WebLogic	7001
Exchange	5985, 5986
Tomcat	1009
JBOSS	9990
Oracle	1521
SQL Server	1433
MySQL	3306

Sample OpenStack SNMP Configuration File for PM Metric Collection

NOTE: This is a sample configuration file for the process found [here \(on page 26\)](#). 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 = UUID: 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%

```