



Workload Optimization Manager 3.6.3 Target Configuration Guide

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any product specific license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/eula>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

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 products and features described in this document as in development or available at a future date remain in varying stages of development and will be offered on a when-and-if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco, and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

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.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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](http://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. (1721R)

© 2023 Cisco Systems, Inc. All rights reserved.

Contents

Target Configuration.....	6
Adding and Removing Targets.....	8
Applications and Databases Targets.....	9
Cisco AppDynamics.....	9
Application Insights.....	15
New Relic.....	17
Dynatrace.....	22
MySQL.....	26
Enabling User Permissions on MySQL Server.....	30
Oracle.....	31
Creating a Service User Account in Oracle.....	36
SQL Server.....	36
JBoss.....	39
Apache Tomcat.....	42
JVM Application.....	45
Oracle WebLogic.....	48
IBM WebSphere.....	51
Instana.....	54
Cloud Native Targets.....	58
Fabric Targets.....	63
Cisco UCS Manager.....	65
HPE OneView.....	67
Guest OS Processes Targets.....	71
WMI.....	71
Creating a WMI User Account.....	74
Enabling WMI.....	74
SNMP.....	75
Enabling SNMP.....	78
Hyperconverged Targets.....	81
Cisco HyperFlex.....	83
Nutanix	85
Pinning Nutanix Controller VMs in Generic Hypervisor Mode.....	89
Hypervisor Targets.....	91
Microsoft Hyper-V.....	94

Creating A Service User Account.....	97
vCenter Server.....	98
Creating A Service User Account In vCenter.....	103
Other Information Imported from vCenter.....	105
Orchestrator Targets.....	106
ActionScript Server.....	106
Flexera One.....	112
ServiceNow.....	116
Private Cloud.....	118
Virtual Machine Manager.....	119
Public Cloud.....	124
Amazon Web Services.....	124
AWS Billing Families.....	133
AWS Billing.....	134
Google Cloud Platform.....	136
Microsoft Azure.....	146
Azure Billing	159
Microsoft Enterprise Agreement.....	162
Azure Enterprise Agreements.....	165
Storage Targets.....	168
EMC VMAX.....	170
EMC XtremIO.....	173
EMC ScaleIO.....	175
EMC VPLEX.....	177
HPE 3PAR.....	178
IBM FlashSystem.....	181
NetApp.....	185
Restricted Service Accounts In NetApp.....	187
Pure Storage	190
Virtual Desktop Infrastructure Targets.....	193
Horizon.....	194
Appendix – Target Configuration.....	197
Cisco Unified Computing System.....	197
Enabling Collection of Memory Statistics: AWS.....	198
Memory Metrics Collection: Azure.....	200
GCP Target Service Account.....	201
GCP Billing Target Service Account.....	205

Enabling Windows Remote Management.....	206
Enabling WinRM Via Global Policy Objects.....	207
Enabling WinRM Via PowerShell.....	208
Secure Setup of WSMAN.....	209
Port Configuration.....	209
AWS Target IAM Role Requirements.....	211



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: `AnalysisException 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.

- Applications and Databases
 - Apache Tomcat 7.x, 8.x, and 8.5.x
 - AppDynamics 4.1+
 - AppInsights
 - Dynatrace 1.1+
 - IBM WebSphere Application Server 8.5+
 - Instana, release-209 or later
 - JBoss Application Server 6.3+
 - JVM 6.0+
 - Microsoft SQL Server 2012, 2014, 2016, 2017, and 2019
 - MySQL 5.6.x and 5.7.x
 - NewRelic
 - Oracle 11g R2, 12c, 18c, and 19c
 - Oracle WebLogic 12c

- Cloud Native
 - Kubernetes, including any compliant k8s distribution (Rancher, Tanzu, open source, etc.)
 - Cloud-hosted k8s services (AKS, EKS, GKE, IBM, Cisco IKS, ROKS, ROSA, etc.)
 - Red Hat OpenShift 3.11 and higher (OCP 4.x)
- Fabric and Network
 - Cisco UCS Manager 3.1+
 - HPE OneView 3.00.04
- Guest OS Processes
 - SNMP
 - WMI: Windows versions 8 / 8.1, 10, 2008 R2, 2012 / 2012 R2, 2016, 2019 and 7
- Hyperconverged
 - Cisco HyperFlex 3.5
 - Nutanix Community Edition
 - VMware vSAN
- Hypervisors
 - Microsoft Hyper-V 2008 R2, Hyper-V 2012/2012 R2, Hyper-V 2016, Hyper-V 2019
 - VMware vCenter 7.0 and 8.0
- Orchestrator
 - ActionScript
 - Flexera One
 - ServiceNow
- Private Cloud
 - Microsoft System Center 2012/2012 R2 Virtual Machine Manager, System Center 2016 Virtual Machine Manager, and System Center Virtual Machine Manager 2019
- Public Cloud
 - Amazon AWS
 - Amazon AWS Billing
 - Google Cloud Platform (GCP)
 - GCP Billing
 - Microsoft Azure Service Principal
 - Azure Billing
 - Microsoft Enterprise Agreement
- Storage
 - EMC ScaleIO 2.x and 3.x
 - EMC VMAX using SMI-S 8.1+
 - EMC VPLEX Local Architecture with 1:1 mapping of virtual volumes and LUNs
 - EMC XtremIO XMS 4.0+
 - HPE 3PAR InForm OS 3.2.2+, 3PAR SMI-S, 3PAR WSAPI
 - IBM FlashSystem running on Spectrum Virtualize 8.3.1.2 or later (8.4.2.0 or later recommended)
 - NetApp Cluster Mode using ONTAP 8.0+ (excluding AFF and SolidFire)
 - Pure Storage F-series and M-series arrays
- 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 TLS 1.2 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**.



Applications and Databases Targets

Applications and Databases 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.

NOTE:

As it manages your applications environment, Workload Optimization Manager discovers connected application components to stitch them into a supply chain of entities. For connections that are made by name and not IP address, Workload Optimization Manager makes DNS calls to resolve these names to IP addresses. This can happen during repeated discovery cycles.

Supply Chain

Applications and Databases 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.

Cisco 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.

NOTE:

For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, Instana, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications. For more information, see [Cloud Native Targets \(on page 58\)](#).

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.

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
Machine (when the machine type is Container)	Container
Server	Virtual 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 or API Client Name@Account

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 or Client Secret

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 @`

- Collect Virtual Machine Metrics

Whether to collect VM metrics from this target. When **Collect Virtual Machine Metrics** is turned on, Workload Optimization Manager uses the VM metrics it collects from this target, instead of the VM data that can optionally be collected from related infrastructure targets.

- Secure Connection

When checked, Workload Optimization Manager will connect via HTTPS. Make sure the required certificate is configured for use on the host.

- Use API Client (OAuth)

When checked, enables Open Authorization (OAuth) token-based authentication for the target connection.

- 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.

- Secure Proxy Connection

When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

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

Actions

NOTE:

The specific actions that Workload Optimization Manager recommends can differ, depending on the processes that Workload Optimization Manager discovers.

For other application components, Workload Optimization Manager can recommend actions based on the resources it can discover for the application. For example, Node.js® applications report CPU usage, so Workload Optimization Manager can generate vCPU resize actions and display them in the user interface.

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

Entity Type	Action
Database Server	<ul style="list-style-type: none"> Resize DB Mem Recommendation only. Resize Connections Recommendation only. Resize Transaction Log Recommendation only.

NOTE:

For different types of Database Servers, the AppDynamics target returns different metrics. This affects Workload Optimization Manager actions as follows:

- MySQL:

For MySQL database servers, analysis does not generate resize actions for DB Memory, Connections, or Transaction Log. The target does not discover DB Cache Hit Rate, DB Memory, Connections, or Transaction Log.
- Microsoft SQL Server:

For Microsoft SQL database servers, analysis does not generate resize actions for DB Memory, Connections, or Transaction Log. The target does not discover DB Memory or Connections.
- MongoDB:

For MongoDB database servers, analysis does not generate resize actions for DB Memory, Connections, or Transaction Log. The target does not discover DB Cache Hit Rate, DB Memory, Transactions, or Transaction Log.
- Oracle:

For Oracle database servers, analysis does not generate resize actions for DB Memory, Connections, or Transaction Log. The target does not discover DB Memory, Connections, or Transaction Log.

Monitored Resources

NOTE:

The exact resources this target monitors can differ based on application type. The following list of metrics per entity includes all resources you might see.

For a given VM, the resources you see depend on how the VM is discovered, and whether the VM provides resources for an application discovered by this target:

- If the VM hosts an application that is discovered through this target, then you will see VM metrics discovered through this target.
- If the VM is discovered through a different target, and it does not host any application discovered through this target, then you will see VM metrics discovered through that different target.
- If the VM is discovered through this target, but it does not host any application discovered through this target, then Workload Optimization Manager does not display metrics for the VM.

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

Entity Type	Commodity
Business Transaction	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Business Applications	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Service	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). For Kubernetes, this is the desired <i>weighted average</i> response time of all Application Component replicas associated with a Service. ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. For Kubernetes, this is the maximum number of transactions per second that each Application Component replica can handle.
Application Component	<ul style="list-style-type: none"> ▪ Virtual CPU (Java, .NET, and Node.js only) Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory (Java, .NET, and Node.js only)

Entity Type	Commodity
	<p>Virtual Memory is the measurement of memory utilized by the entity.</p> <ul style="list-style-type: none"> ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Heap (Java, .NET, and Node.js only) Heap is the portion of a VM or container’s memory allocated to individual applications. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Connections Connection is the measurement of Database Server connections utilized by applications. ▪ Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC). ▪ Threads Threads is the measurement of thread capacity utilized by applications.
Database Server	<ul style="list-style-type: none"> ▪ Transactions (SQL, MySQL, and Oracle only) Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Connections (MongoDB only) Connection is the measurement of Database Server connections utilized by applications. ▪ Transaction Log (SQL only) The utilization of the server’s capacity for storage devoted to transaction logs ▪ DB Cache Hit Rate (SQL and Oracle only) DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. <p style="text-align: center;">NOTE: To collect data, a machine agent must be present and database hardware monitoring must be enabled.</p> <ul style="list-style-type: none"> ▪ Virtual Memory

Entity Type	Commodity
	Virtual Memory is the measurement of memory utilized by the entity. NOTE: To collect data, a machine agent must be present and database hardware monitoring must be enabled.

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 \(on page 146\)](#) 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 ID**
The Client 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.
- 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.
- Secure Proxy Connection
 - When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

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

NOTE:

The specific actions that Workload Optimization Manager recommends can differ, depending on the processes that Workload Optimization Manager discovers.

For other application components, Workload Optimization Manager can recommend actions based on the resources it can discover for the application. For example, Node.js® applications report CPU usage, so Workload Optimization Manager can generate vCPU resize actions and display them in the user interface.

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 Virtual CPU is the measurement of CPU utilized by the entity. ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms).

New Relic

Workload Optimization Manager supports workload management of the application infrastructure monitored by New Relic, from application instance to host. With information obtained from New Relic, Workload Optimization Manager can make recommendations and take actions to both assure performance and drive efficiency to address the demands of each individual application. For Kubernetes environments, Workload Optimization Manager stitches containerized application components into the supply chain to provide a unified view of your applications.

NOTE:

For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, Instana, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications. For more information, see [Cloud Native Targets \(on page 58\)](#).

Prerequisites

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

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 Term	Workload Optimization Manager Term
APM: Key Transactions	Business Transaction
APM: Application / Service (New Relic One)	Service
APM: Application Instance	Application Component
Infra: Database	Database Server
Infra: Host	Virtual Machine

For VM entities

Supported Applications

Workload Optimization Manager discovers the following application types (and associated commodities) via the New Relic target:

Application Type	Commodities
.NET	Virtual CPU, Virtual Memory, Response Time, Transactions
GO	Virtual CPU, Virtual Memory, Response Time, Transactions
Java	Virtual CPU, Virtual Memory, Response Time, Transactions, Heap, Collection Time, Threads
Node.js	Virtual CPU, Virtual Memory, Response Time, Transactions, Heap, Collection Time
PHP	Virtual CPU, Virtual Memory, Response Time, Transactions
Python	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
MS SQL	Cache Hit Rate, Virtual Memory, Transactions
MySQL	Cache Hit Rate NOTE: Workload Optimization Manager no longer shows DB Cache Hit Rate values since it is not supported in New Relic MySQL version 8.0 and higher. For more information, see the New Relic documentation .
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 [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 GraphiQL Explorer](#).
- EU Region
If checked, Workload Optimization Manager will use the EU API endpoints.
- Collect Virtual Machine Metrics
Whether to collect VM metrics from this target. When **Collect Virtual Machine Metrics** is turned on, Workload Optimization Manager uses the VM metrics it collects from this target, instead of the VM data that can optionally be collected from related infrastructure targets.
- 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.
- Secure Proxy Connection
When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Actions

NOTE:

The specific actions that Workload Optimization Manager recommends can differ, depending on the processes that Workload Optimization Manager discovers.

For other application components, Workload Optimization Manager can recommend actions based on the resources it can discover for the application. For example, Node.js® applications report CPU usage, so Workload Optimization Manager can generate vCPU resize actions and display them in the user interface.

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 this target monitors can differ based on application type. The following list of metrics per entity includes all resources you might see.

For a given VM, the resources you see depend on how the VM is discovered, and whether the VM provides resources for an application discovered by this target:

- If the VM hosts an application that is discovered through this target, then you will see VM metrics discovered through this target.
- If the VM is discovered through a different target, and it does not host any application discovered through this target, then you will see VM metrics discovered through that different target.
- If the VM is discovered through this target, but it does not host any application discovered through this target, then Workload Optimization Manager does not display metrics for the VM.

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 Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity.

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="586 281 769 306">▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. <li data-bbox="586 390 686 415">▪ Heap Heap is the portion of a VM or container's memory allocated to individual applications. <li data-bbox="586 499 800 525">▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). <li data-bbox="586 636 756 661">▪ Connection Connection is the measurement of Database Server connections utilized by applications. <li data-bbox="586 745 889 770">▪ Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC). <li data-bbox="586 854 719 879">▪ Threads Threads is the measurement of thread capacity utilized by applications.
Database	<ul style="list-style-type: none"> <li data-bbox="586 968 1385 1035">▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. <li data-bbox="586 1052 1479 1144">▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. <li data-bbox="586 1161 1466 1253">▪ Database memory Database memory (or DBMem) is the measurement of memory utilized by a Database Server. <li data-bbox="586 1270 1466 1362">▪ Connections Connection is the measurement of Database Server connections utilized by applications. <li data-bbox="586 1379 1498 1501">▪ DB Cache Hit Rate DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency.
Business Transaction	<ul style="list-style-type: none"> <li data-bbox="586 1545 1482 1667">▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms).

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Service	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity.

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.

NOTE:

For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, Instana, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications. For more information, see [Cloud Native Targets \(on page 58\)](#).

Prerequisites

- A Dynatrace Server instance
This instance must be configured to monitor applications running in your environment.
Workload Optimization Manager supports both SaaS and on-prem Dynatrace server installations.
- Managed VMs that host 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.
- An API access token with the proper scopes

Workload Optimization Manager uses the API token to authenticate its calls to the Dynatrace API. This token must have permission to execute GET methods via the Dynatrace API, both Version 1 and Version 2. Generate a new generic access token with these scopes:

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> - API V1 scopes: <ul style="list-style-type: none"> • Access problem and event feed, metrics, and topology - API V2 scopes: <ul style="list-style-type: none"> • Read entities • Read metrics

NOTE:

If you are updating to Workload Optimization Manager version 3.4.2 or later, from a version that is earlier than 3.4.2, you must generate a new API token for each existing Dynatrace target. Then you must enter that token in the target configuration, and validate the target.

If the target still fails to validate after you update the access token, take note of your configuration settings, delete the target, and configure the target again. Be sure to use the new API token that you have generated.

Entity Mapping

After validating the new target, Workload Optimization Manager discovers the connected entities. The entity names that Workload Optimization Manager displays in the Supply Chain differ from the entity names that Dynatrace displays in its user interface, as follows:

Dynatrace Naming	Workload Optimization Manager Entity
Application	Business Application NOTE: For Dynatrace Applications, Workload Optimization Manager displays Business Application entities in the supply chain when they have been active for at least three days.
Service	Service
Process	Application Component, Database Server
NA	Container
Host	Virtual Machine

Adding a Dynatrace Target

NOTE:

You can manage 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 manage 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**
 For an on-prem installation of Dynatrace, give the host name or IP and endpoint, separated by a slash. For example, `10.10.10.10/e/b70e3eb2-e82b-4c13-a5a4-560d9865841r`.
 For a SaaS installation, provide the URL *without* the `{IP_ADDRESS}/e/` statement in the path. For example, `ayz12745.live.dynatrace.com`
- **API Token**
 The token that Workload Optimization Manager can use to authenticate its calls to the Dynatrace API. This token must have permission to execute GET methods via the Dynatrace API V1 and V2. Refer to the Prerequisites section for more information.
- **Collect Virtual Machine Metrics**
 Whether to collect VM metrics from this target. When **Collect Virtual Machine Metrics** is turned on, Workload Optimization Manager uses the VM metrics it collects from this target, instead of the VM data that can optionally be collected from related infrastructure targets.
- **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.
- **Secure Proxy Connection**
 Whether to use a secure connection with the proxy. When this is turned on, Workload Optimization Manager connects to the proxy through HTTPS.

Actions

Entity:	Actions:
Application Component	<ul style="list-style-type: none"> ▪ Resize <ul style="list-style-type: none"> – Heap Recommended, only.

Monitored Resources

NOTE:

The subset of resources that Workload Optimization Manager discovers for an application depends on the application type. The following list of metrics per entity includes the full set of resources Workload Optimization Manager can discover for Dynatrace applications.

For Database Server applications, Workload Optimization Manager only discovers metrics for MySQL and MSSQL databases.

For a given VM, the resources you see depend on how the VM is discovered, and whether the VM provides resources for an application discovered by this target:

- If the VM hosts an application that is discovered through this target, then you will see VM metrics discovered through this target.
- If the VM is discovered through a different target, and it does not host any application discovered through this target, then you will see VM metrics discovered through that different target.
- If the VM is discovered through this target, but it does not host any application discovered through this target, then Workload Optimization Manager does not display metrics for the VM.

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 Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Service	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Application Component	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Remaining GC Capacity (Java only) Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC).

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Heap (Java only) Heap is the portion of a VM or container’s memory allocated to individual applications.
Database Server	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Database memory (SQL and MySQL only) Database memory (or DBMem) is the measurement of memory utilized by a Database Server. Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. ▪ DB Cache Hit Rate (SQL only) DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency. ▪ Transactions (SQL only) Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Container	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity.

MySQL

To manage MySQL databases, Workload Optimization Manager can connect to one or more database servers within a defined scope.

Prerequisites

- User Permissions are enabled on the MySQL Server. See [Enabling User Permissions on MySQL \(on page 30\)](#)

Adding a MySQL Database Target

You can add all matching targets within a given scope.

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

- **Target ID**
Name displayed in the Workload Optimization Manager UI
- **Username**
Username of the account Workload Optimization Manager uses to connect to the target.
- **Password**
Password of the account Workload Optimization Manager uses to connect to the target.
- **Scope**
A group of applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.
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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.
- **Port Number**
The MySQL remote port. If blank, Workload Optimization Manager will use the MySQL default port of 3306.
- **Full Validation**
When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

Actions

Entity:	Actions:
Database Server	<ul style="list-style-type: none"> ▪ Resize (Recommend, only) <ul style="list-style-type: none"> – Database memory (DBMem) Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. Workload Optimization Manager uses database memory and cache hit rate data to decide whether resize actions are necessary. A high cache hit rate value indicates efficiency. The optimal value is 100% for on-prem (self-hosted) Database Servers, and 90% for cloud Database Servers. When the cache hit rate reaches the optimal value, no action generates even if database memory utilization is high. If utilization is low, a resize down action generates. When the cache hit rate is below the optimal value but database memory utilization remains low, no action generates. If utilization is high, a resize up action generates.

Entity:	Actions:
	<ul style="list-style-type: none"> - Connections <p>Workload Optimization Manager uses connection data to generate memory resize actions for on-prem Database Servers.</p>
Virtual Machine	<ul style="list-style-type: none"> ▪ Resize <ul style="list-style-type: none"> - Resize resource capacity <p>Change the capacity of a resource that is allocated for the VM. For example, a resize action might recommend increasing the VMem available to a VM. Before recommending this action, Workload Optimization Manager verifies that the VM's cluster can adequately support the new size. If the cluster is highly utilized, Workload Optimization Manager will recommend a move action, taking into consideration the capacity of the new cluster and compliance with existing placement policies.</p> <p>For hypervisor targets, Workload Optimization Manager can resize vCPU by changing the VM's socket or cores per socket count. For details, see "vCPU Scaling Controls" in the <i>User Guide</i>.</p> - Resize resource reservation <p>Change the amount of a resource that is reserved for a VM. For example, a VM could have an excess amount of memory reserved. That can cause memory congestion on the host – A resize action might recommend reducing the amount reserved, freeing up that resource and reducing congestion</p> - Resize resource limit <p>Change the limit that is set on the VM for a resource. For example, a VM could have a memory limit set on it. If the VM is experiencing memory shortage, an action that decreases or removes the limit could improve performance on that VM.</p> ▪ Move <p>Move a VM due to:</p> <ul style="list-style-type: none"> - High resource utilization on VM or host - Excess IOPS or latency in VStorage - Workload placement violation - Underutilized host (move VM before suspending host) ▪ Reconfigure <p>Change a VM's configuration to comply with a policy.</p>

Entity:	Actions:
	For hypervisor targets, Workload Optimization Manager can reconfigure VMs that violate vCPU scaling policies. For details, see "vCPU Scaling Controls" in the <i>User Guide</i> .

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="586 667 1524 840"> ▪ Database memory Database memory (or DBMem) is the measurement of memory utilized by a Database Server. Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. <li data-bbox="586 850 1524 945"> ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. <li data-bbox="586 955 1524 1081"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). <li data-bbox="586 1092 1524 1186"> ▪ Connections Connection is the measurement of Database Server connections utilized by applications. <li data-bbox="586 1197 1524 1323"> ▪ DB Cache Hit Rate DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency.
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="586 1373 1524 1438"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. <li data-bbox="586 1449 1524 1514"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. <li data-bbox="586 1524 1524 1589"> ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. <li data-bbox="586 1600 1524 1665"> ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. <li data-bbox="586 1675 1524 1761"> ▪ Latency Latency is the measurement of storage latency utilized by the entity.

Enabling User Permissions on MySQL Server

Follow the following steps in order to enable appropriate user permissions on a MySQL Server.

1. Edit the MySQL server's configuration file.

You must edit the `.conf` file on the MySQL server to grant user permissions. Open a secure shell session on the server and edit the file. Depending on the platform your MySQL is running on, you'll find the file at different locations:

- Debian Linux:
`/etc/mysql/my.cnf`
- Red Hat Linux (Fedora or Centos):
`/etc/my.cnf`
- FreeBSD Linux:
You must create the file at `/var/db/mysql/my.cnf`

Open the file in an editor and find the section, `[mysqld]`. Then make the following changes:

- Comment out the line:
`skip-networking`
Commenting out this line enables remote connections over TCP/Is.
- Bind your MySQL server address
In the config file, add the line:
`bind-address=<MySQL_IP_Address>`
- Enable the collection of Transaction metrics
In the config file, add the line:
`innodb_monitor_enable = trx_rw_commits, trx_nl_ro_commits, trx_ro_commits, trx_rollbacks`

For example, if your MySQL server has the address, `123.45.66.77`, then after you have bound the IP address and enabled Transaction metrics, the section of the `.conf` file should appear as follows:

```
[mysqld]

user                    = mysql

pid-file                = /var/run/mysqld/mysqld.pid

socket                  = /var/run/mysqld/mysqld.sock

port                    = 3306

basedir                 = /usr

datadir                 = /var/lib/mysql

tmpdir                  = /tmp

language                 = /usr/share/mysql/English

bind-address            = 123.45.66.77

# skip-networking
```

```
# Uncomment the following line for MySQL versions 5.6+
innodb_monitor_enable = trx_rw_commits, trx_nl_ro_commits, trx_ro_commits, trx_rollbacks
```

```
....
```

When you are done, save the `.conf` file.

NOTE: Some MySQL installations use multiple configuration files. If a setting you made does not have the desired effect, make sure that a different configuration file is not overwriting the value.

2. Enable collection of Response Time metrics.

Execute the following command to log into to the MySQL server:

```
$mysql -u root -p mysql
```

Then execute the following SQL commands:

```
UPDATE performance_schema.setup_instruments SET ENABLED = 'YES' WHERE NAME LIKE 'statement/sql%';
```

```
UPDATE performance_schema.setup_instruments SET TIMED = 'YES' WHERE NAME LIKE 'statement/sql%';
```

NOTE:

If you want these changes to take effect each time you restart the MySQL server, add these statements to a file, and start the server with the `--init-file` option. For example, if you name the file `MyInit.txt`, then start the MySQL server with the following option:

```
--init-file=MyInit.txt
```

3. Give your Workload Optimization Manager server remote access to the database.

If you are not already logged into the MySQL server, execute the following command:

```
$mysql -u root -p mysql
```

Then execute the following commands:

Assume a user named `USER_NAME` with a password `PWD_STRING`. Then assume that your Workload Optimization Manager has an IP address of `10.10.123.45`. The following command grants privileges to that Workload Optimization Manager, if it connects with the specified user account:

```
GRANT SELECT ON performance_schema.* TO 'USER_NAME'@'10.10.123.45' IDENTIFIED BY 'PWD_STRING';
GRANT PROCESS ON *.* TO 'USER_NAME'@'10.10.123.45' IDENTIFIED BY 'PWD_STRING';
FLUSH PRIVILEGES;
```

Note that the `FLUSH PRIVILEGES` command causes MySQL to retain these settings upon restart.

When you're finished running these SQL commands, log out of MySQL.

Oracle

To connect to an Oracle database, you will:

- Add a Dynamic Performance view to the Oracle database
- Configure a service account on the database that Workload Optimization Manager can use to log on
- Find the Service Name and port for the database

Prerequisites

- User permissions that grant access to Workload Optimization Manager through a specific user account. See [Creating a Service User Account in Oracle \(on page 36\)](#).
- Dynamic Performance View (V\$) must be enabled. See "Adding a Dynamic Performance View", below.
- Access through the firewall to the Oracle database port that you specify for the Workload Optimization Manager target connection

Adding a Dynamic Performance View

In order to collect data from the Oracle database, Workload Optimization Manager uses the Dynamic Performance View (referred to as V\$). V\$ is not enabled by default. You must run a script to build the tables and views that are necessary to enable V\$. In some environments only the DBA has privileges to run this script.

To enable V\$:

- Open a secure shell session (ssh) on the database host as a system user or a user with the `sysdba` role
- In the shell session enter the following commands:

```
sqlplus /nolog
connect /as sysdba
CREATE USER My_Username IDENTIFIED BY My_Password container=all;
GRANT CONNECT TO My_Username container=all;
GRANT sysdba TO My_Username container=all;
```

NOTE:

If security or other practices prohibit assigning SYSDBA to this user, you can use the following command to provide access to all V\$ views:

```
GRANT select any dictionary TO My_Username;
```

This creates a user account named `My_Username` with full privileges to access the V\$ Dynamic Performance view.

Adding an Oracle Database to Workload Optimization Manager

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

To add a database server as a target, specify:

- **Target Name**
The target name that will display in the Workload Optimization Manager user interface.
- **Username/Password**
Valid client credentials for the database server. For Workload Optimization Manager to execute actions, the account must have administrator privileges. Also, you must have enabled user permissions to this user account, including remote access from the Workload Optimization Manager server.
- **Scope**
A group of applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.
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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.

NOTE:

All database servers in the scope must share the same service name, credentials, and port. For databases that have a different value for any of these, you must create a separate target using those values.

- Oracle Port
The port that connects to the database. You must open the firewall on the database server to allow access through this port. For further information, see "Finding the Service Name and Port", below.
- Oracle Service Name
The service name for the database that you are connecting to.
- Full Validation
When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

Finding the Service Name and Port

To specify an Oracle target, you must provide the service name and port that you want to connect to. To find the service name for your database:

- Open a secure shell session (ssh) on the database host as a system user or a user with the `sysdba` role
- In the shell session, enter the command, `lsnrctl status`
Find the line that has the string `PROTOCOL=TCP` and note the port number.
- In the shell session enter the following commands:

```
sqlplus /nolog
connect /as sysdba
SELECT SYS_CONTEXT('userenv', 'db_name') FROM dual;
```

Note the service name that displays as a result of these commands.

Actions

Entity:	Actions:
Database Server	<p>Resize (Recommend, only)</p> <ul style="list-style-type: none"> ▪ Connections Workload Optimization Manager uses connection data to generate memory resize actions for on-prem Database Servers. ▪ Database memory (DBMem) Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. Workload Optimization Manager uses database memory and cache hit rate data to decide whether resize actions are necessary. A high cache hit rate value indicates efficiency. The optimal value is 100% for on-prem (self-hosted) Database Servers, and 90% for cloud Database Servers. When the cache hit rate reaches the optimal value, no action generates even if database memory utilization is high. If utilization is low, a resize down action generates.

Entity:	Actions:
	<p>When the cache hit rate is below the optimal value but database memory utilization remains low, no action generates. If utilization is high, a resize up action generates.</p>
Virtual Machine	<ul style="list-style-type: none"> ▪ Resize <ul style="list-style-type: none"> – Resize resource capacity <p>Change the capacity of a resource that is allocated for the VM. For example, a resize action might recommend increasing the VMem available to a VM. Before recommending this action, Workload Optimization Manager verifies that the VM's cluster can adequately support the new size. If the cluster is highly utilized, Workload Optimization Manager will recommend a move action, taking into consideration the capacity of the new cluster and compliance with existing placement policies.</p> <p>For hypervisor targets, Workload Optimization Manager can resize vCPU by changing the VM's socket or cores per socket count. For details, see "VCPU Scaling Controls" in the <i>User Guide</i>.</p> – Resize resource reservation <p>Change the amount of a resource that is reserved for a VM. For example, a VM could have an excess amount of memory reserved. That can cause memory congestion on the host – A resize action might recommend reducing the amount reserved, freeing up that resource and reducing congestion</p> – Resize resource limit <p>Change the limit that is set on the VM for a resource. For example, a VM could have a memory limit set on it. If the VM is experiencing memory shortage, an action that decreases or removes the limit could improve performance on that VM.</p> ▪ Move <p>Move a VM due to:</p> <ul style="list-style-type: none"> – High resource utilization on VM or host – Excess IOPS or latency in VStorage – Workload placement violation – Underutilized host (move VM before suspending host) ▪ Reconfigure <p>Change a VM's configuration to comply with a policy.</p>

Entity:	Actions:
	For hypervisor targets, Workload Optimization Manager can reconfigure VMs that violate vCPU scaling policies. For details, see "vCPU Scaling Controls" in the <i>User Guide</i> .

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="586 667 1524 835"> ▪ Database memory Database memory (or DBMem) is the measurement of memory utilized by a Database Server. Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. <li data-bbox="586 846 1524 947"> ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. <li data-bbox="586 957 1524 1087"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). <li data-bbox="586 1098 1524 1192"> ▪ Connections Connection is the measurement of Database Server connections utilized by applications. <li data-bbox="586 1203 1524 1276"> ▪ TransactionLog The utilization of the server's capacity for storage devoted to transaction logs <li data-bbox="586 1287 1524 1409"> ▪ DB Cache Hit Rate DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency.
Virtual Machine	<ul style="list-style-type: none"> <li data-bbox="586 1453 1524 1526"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. <li data-bbox="586 1537 1524 1610"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. <li data-bbox="586 1621 1524 1694"> ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. <li data-bbox="586 1705 1524 1778"> ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. <li data-bbox="586 1789 1524 1862"> ▪ Latency Latency is the measurement of storage latency utilized by the entity.

Creating a Service User Account in Oracle

To collect data from the Oracle database, Workload Optimization Manager requires a service account that has privileges to access the V\$ Dynamic Performance view. To create this account:

- Open a secure shell session (ssh) on the database host as a system user or a user with the `sysdba` role
- In the shell session enter the following commands:

```
sqlplus /nolog
connect /as sysdba
CREATE USER My_Username IDENTIFIED BY My_Password container=all;
GRANT CONNECT TO My_Username container=all;
GRANT sysdba TO My_Username container=all;
```

This creates a user account named `My_Username` with full privileges to access the V\$ Dynamic Performance view.

NOTE:

The above example uses a fictitious username. To comply with Oracle 12C norms, the username should include a prefix of `c##`.

Some enterprises don't allow accounts with `sysdba` access. Cisco recommends using `sysdba`, according to the Oracle documentation. However, you can work with your Oracle DBA staff to provide read access to the following views, which are the ones that Workload Optimization Manager needs:

- V\$INSTANCE
- V\$LOG
- V\$LOGFILE
- V\$PARAMETER
- V\$PGASTAT
- V\$RESOURCE_LIMIT
- V\$SGASTAT
- V\$SYS_TIME_MODEL
- V\$SYSMETRIC
- V\$SYSSTAT

SQL Server

Workload Optimization Manager supports the following versions of this target:

Microsoft SQL Server 2012, 2014, 2016, 2017, and 2019

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.

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
- AD Domain
The Active Directory domain used by Workload Optimization Manager in conjunction with the Username for authentication. Leave blank for local accounts.
- Password
Password for the account. This username must not include the AD domain
- Scope
A group of applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.

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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.
- Browsing Service Port
The UDP port for the browsing service that listens for incoming connections to the SQL instances running on the SQL Server. If the browsing service is reachable via the specified port, Workload Optimization Manager discovers the SQL instances used by the VM group that you defined as your scope, as well as the listening ports on those SQL instances. The default UDP port is 1434.

If the service is unreachable, or if you did not specify a UDP port, Workload Optimization Manager will use the TCP port that you specified in the SQLServer Port field to discover SQL instances. See the next item.
- SQLServer Port

The TCP port for the SQL Server. Workload Optimization Manager will use this port if the browsing service is unreachable, or if you did not specify a browsing service port. The default TCP port is 1433.

- Full Validation

When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server 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	Without separate targets to discover Guest OS Processes or Application Servers, Workload Optimization Manager does not generate actions on applications. Instead, it generates resize actions on the host VMs. For on-prem environments, if host utilization is high enough on the physical machine running the application VM, Workload Optimization Manager can also recommend provisioning a new host.
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 ▪ Suspend VM ▪ Provision VM

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"> ▪ Database memory Database memory (or DBMem) is the measurement of memory utilized by a Database Server. Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Connections Connection is the measurement of Database Server connections utilized by applications. ▪ TransactionLog

Entity Type	Commodity
	<p>The utilization of the server's capacity for storage devoted to transaction logs</p> <ul style="list-style-type: none"> ▪ DB Cache Hit Rate <p>DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency.</p>
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. ▪ Latency Latency is the measurement of storage latency utilized by the entity.

JBoss

Workload Optimization Manager supports connecting to JBoss targets running in these operation modes:

- **Managed Domain**
A collection of JBoss servers in a domain, with a single Domain Controller process acting as the central management control point. In this case, the VM that hosts the Domain Controller will be the target. You configure a Domain Controller via the `domain-controller` entry in the `host.xml` file on the JBoss machine.
- **Standalone**
A single JBoss server. The VM that hosts the server is the target.

Prerequisites

- The target VM is properly configured as Domain Controller or standalone, depending on the operation mode.
- Workload Optimization Manager discovers JBoss servers that are running on VMs, or on containers that are deployed on VMs. To set the target for a JBoss server, you must have first discovered the hosting VM through a hypervisor target. For information about hypervisor targets, see [Hypervisor Targets \(on page 91\)](#).
- Workload Optimization Manager monitors the Threads resource in application servers to track utilization of thread pool capacity. To monitor threads in JBoss, each JBoss server must define a thread pool in its configuration files. For instructions on specifying thread pools, see the JBoss documentation.
- JBoss supports Role-Based Access Control (RBAC) as a way to specify a set of permissions for user accounts. JBoss disables RBAC by default. If you enable RBAC, Workload Optimization Manager requires the `Monitor` role for the target JBoss account in order to discover JBoss servers.

Adding a JBoss Target

You can add an individual JBoss 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 user interface
- **Username**
The username for a JBoss account that has a role with the necessary permissions.
- **Password**
The password for the JBoss account that you will use.
- **Scope**
A group of JBoss servers, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.
After 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.
- **JBoss Webservice port**
The port for connections to the JBoss server. The default port for HTTP access is 9990.
- **Full Validation**
When enabled, Workload Optimization Manager will require all JBoss servers hosted on the VMs in the selected scope to be valid targets. If Workload Optimization Manager is unable to authenticate a JBoss server in the scope, the target will not validate and data will not be collected.
- **Secure Connection**
When enabled, Workload Optimization Manager will connect to JBoss servers via HTTPS. Be sure that the required certificate has been configured for use on the host.
- **Proxy Information**
Specify proxy information only if you connect to a JBoss instance via proxy.
 - **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.
 - **Secure Proxy Connection**
When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Application Component Names

Workload Optimization Manager displays discovered JBoss servers in the user interface. These names indicate whether they're standalone or in a managed domain.

The name is divided into three sections:

- Domain Controller name or “STANDALONE”
- The server name
- The name or IP address of the VM hosting the JBoss server

For example:

STANDALONE:acm-jboss73 [ACM-JBoss7.3-171.47]

Actions

Entity Type	Action
Application Component	<ul style="list-style-type: none"> ▪ Resize Heap Recommendation only. ▪ Resize Thread Pool Recommendation only. ▪ Resize Connection Capacity Recommendation only.

Monitored Resources

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

Entity Type	Commodity
Application Component	<ul style="list-style-type: none"> ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Heap Heap is the portion of a VM or container’s memory allocated to individual applications. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Threads Threads is the measurement of thread capacity utilized by applications. ▪ Connection Connection is the measurement of Database Server connections utilized by applications. ▪ Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC).

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 JVM 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 \(on page 58\)](#)
- For information about hypervisor targets, see [Hypervisor Targets \(on page 91\)](#)

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 applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.
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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.
- JMX Remote Port**
A JMX port that is set to an already running Tomcat process.
- Full Validation**

When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

- 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.
- Secure Proxy Connection
When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Actions

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

Entity Type	Action
Application Component (Tomcat Application)	<ul style="list-style-type: none"> ▪ Resize Heap Recommendation only. ▪ Resize Thread Pool Recommendation only. ▪ Resize Connection Capacity 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 ▪ Suspend VM ▪ Provision VM

Monitored Resources

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

Entity Type	Commodity
Application Component	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Heap Heap is the portion of a VM or container's memory allocated to individual applications. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Threads Threads is the measurement of thread capacity utilized by applications. ▪ Connection Connection is the measurement of Database Server connections utilized by applications. ▪ Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC).
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.

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 JVM version 6.0 or higher
- 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 \(on page 58\)](#)
- For information about hypervisor targets, see [Hypervisor Targets \(on page 91\)](#)

Configuring JMX Remote Access

Workload Optimization Manager monitors and controls JVM applications 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 JVM Application Targets

When you configure JVM targets, you declare a given scope and add all matching applications within that given scope. To do this, specify:

- **Scope:**

A group of applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.

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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.

- **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: VMs_myCluster.mycorp.com
 - Port number: 123
 - Username: AppUser
 - Password: *****

Then click **ADD**.

- Specify the second target with the following parameters:
 - Scope: 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
Application Component (JVM Application)	<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 ▪ Suspend VM ▪ Provision VM

Monitored Resources

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

Entity Type	Commodity
Application Component (JVM Application)	<ul style="list-style-type: none"> ▪ Heap Heap is the portion of a VM or container's memory allocated to individual applications. ▪ Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC). Data is collected if JVM profiler is enabled.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.

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 \(on page 58\)](#)
- For information about hypervisor targets, see [Hypervisor Targets \(on page 91\)](#)

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 applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.

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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.
- **Port Number**
The WebLogic remote port
- **Full Validation**
When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server 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

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

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 Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Application Component	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Heap Heap is the portion of a VM or container’s memory allocated to individual applications. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Threads Threads is the measurement of thread capacity utilized by applications. ▪ Connection Connection is the measurement of Database Server connections utilized by applications.

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC).
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.

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 WebSphere installation, you can use the WebSphere 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 \(on page 58\)](#)
- For information about container targets, see [Hypervisor Targets \(on page 91\)](#)

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 applications, stitched to the underlying VMs when the VMs are discovered as part of a separate Workload Optimization Manager target.
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 can result in poor performance for discovery and monitoring. To target a larger number of VMs by scope, you should split them across smaller groups and set each group as the scope for a separate target.
- **Port Number**
The WebSphere remote port
- **Full Validation**
When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server 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

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

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 Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Application Component	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity. ▪ Heap Heap is the portion of a VM or container's memory allocated to individual applications. ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Threads Threads is the measurement of thread capacity utilized by applications. ▪ Connection Connection is the measurement of Database Server connections utilized by applications.

Entity Type	Commodity
	<ul style="list-style-type: none"> Remaining GC Capacity Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC).
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.

Instana

Workload Optimization Manager supports discovery of applications that are managed by the Instana platform. Workload Optimization Manager includes the discovered information about these applications in its calculations for environment health.

NOTE:

For Kubernetes environments, Workload Optimization Manager stitches NewRelic, AppDynamics, Instana, and Dynatrace containerized application components into the supply chain to provide a unified view of your applications. For more information, see [Cloud Native Targets \(on page 58\)](#).

Prerequisites

- Instana, release-209 or later
- A self-hosted or SaaS Instana instance
This instance must be configured to monitor applications running in your environment.
- Virtual machines and/or containers that are running applications managed by Instana
For Workload Optimization Manager to discover applications through Instana, the applications must be running on virtual machines or in Kubernetes containers located in your environment. Also, these entities must be managed by Workload Optimization Manager hypervisor, public cloud, or Kubernetes targets.
- A valid Instana API token
This must be an API token created with the default permissions.

Entity Mapping

After validating the new target, Workload Optimization Manager discovers the connected entities, and adds them to the supply chain. This table maps terms from Instana to the types of entities Workload Optimization Manager creates in the supply chain:

Instana Term	Workload Optimization Manager Entity
Application	Business Application
Database	Database Server
Endpoint	Business Transaction
Service	Service
Process	Application Component
Docker / Crio Container	Container

Instana Term	Workload Optimization Manager Entity
Host	Virtual Machine

Instana Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> Default Workload Optimization Manager uses the default API Token for monitoring access.

Adding an Instana Target

NOTE:

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

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

To add an Instana server instance as a target, specify:

- Hostname or IP Address
 The host name or IP of the Instana server.
- API Token
 The Instana API Token.
- Collect Virtual Machine Metrics
 Whether to collect VM metrics from this target. When **Collect Virtual Machine Metrics** is turned on, Workload Optimization Manager uses the VM metrics it collects from this target, instead of the VM data that can optionally be collected from related infrastructure targets.
- Proxy Host
 The address of the proxy used for this target. Only fill out proxy information if you connect to the Instana 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.
- Secure Proxy Connection
 When checked, Workload Optimization Manager will connect to the proxy over HTTPS.

Actions

Workload Optimization Manager does not recommend actions for Instana business applications, business transactions, or services. However, it considers resource utilization by these applications when recommending actions for the underlying virtual machines or containers. Workload Optimization Manager propagates this underlying information upward in the supply chain to show the impact infrastructure risks can have on the business applications, transactions, and services.

In addition, Workload Optimization Manager recommends actions for:

Entity Type	Action
Application Component	<ul style="list-style-type: none"> ▪ Resize Heap 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 Instana supply chain:

Entity Type	Commodity
Business Application	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Business Transaction	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Service	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Transactions Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.
Application Component	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Remaining GC Capacity (Java only) Remaining GC capacity is the measurement of Application Component uptime that is <i>not</i> spent on garbage collection (GC).

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Heap (Java only) Heap is the portion of a VM or container's memory allocated to individual applications.
Database Server	<ul style="list-style-type: none"> ▪ Virtual CPU (Oracle and MySQL) Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory (Oracle and MySQL) Virtual Memory is the measurement of memory utilized by the entity. ▪ Database memory (Oracle and MySQL) Database memory (or DBMem) is the measurement of memory utilized by a Database Server. Actions to resize database memory are driven by data on the Database Server, which is more accurate than data on the hosting VM. ▪ DB Cache Hit Rate (Oracle and MySQL) DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency. ▪ Transaction Log (Oracle) Transaction log is the measurement of storage capacity utilized by Database Servers for transaction logging. ▪ Connections (Oracle) Connection is the measurement of Database Server connections utilized by applications.
Container	<ul style="list-style-type: none"> ▪ Virtual Memory (Docker containers only) Virtual Memory is the measurement of memory utilized by the entity.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity.



Cloud Native Targets

Containers support separation of concerns in a way similar to virtual machines, but allow greater flexibility and use far less overhead. Containers can 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.

To support cloud native environments, Workload Optimization Manager targets Kubernetes clusters. Workload Optimization Manager supports target clusters managed on Kubernetes v1.8 or higher, whether the clusters are managed directly via kubeadm, or via other platforms including:

- Red Hat OpenShift
- Pivotal Kubernetes Service
- Amazon Elastic Kubernetes Service (EKS)
- Azure Kubernetes Service (AKS)
- Cisco Container Platform (CCP)
- Google Kubernetes Engine (GKE)

With Cloud Native targets, Workload Optimization Manager discovers entities related to container platforms in your environment. Discovery can also stitch the container cluster entities together with managed applications. For example, discovery can show the full application stack if your container environment includes applications managed by the following technologies, and you have added them as targets to Workload Optimization Manager:

- [Cisco AppDynamics \(on page 9\)](#)
- [Instana \(on page 54\)](#)
- [Dynatrace \(on page 22\)](#)
- [New Relic \(on page 17\)](#)

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	<p>None</p> <p>No actions are recommended at this level of the supply chain. Instead, actions that affect the service are generated and executed on underlying entities.</p>
Application Component	<p>Suspend</p> <p>Application components are suspended due to a node (virtual machine) suspension</p> <p>APM Actions</p> <p>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.</p>
Container	<p>Resize Container Up/Down</p> <p>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</p> <p>Containers are suspended due to a node (virtual machine) suspension</p>
Container Pod	<p>Move Pod</p> <p>Pods will be moved across nodes (Virtual Machines).</p> <p>Suspend</p> <p>Container Pods are suspended due to a node (virtual machine) suspension</p>
Container Spec	<p>None</p> <p>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</p> <p>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</p> <p>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</p> <p>The following resources may be provisioned:</p> <ul style="list-style-type: none"> ▪ VMem ▪ VCPU ▪ VMem Requests ▪ VCPU Requests ▪ Number of Consumers <p>Suspend</p>

Entity	Action		
	<p>Nodes (virtual machines) may be suspended.</p> <p>Infrastructure-dependent Actions</p> <p>Depending on the technology the node is stitched to for underlying infrastructure, there may be additional actions:</p> <table border="1" data-bbox="618 436 1433 600"> <tr> <td data-bbox="618 436 862 600">On-prem VMware:</td> <td data-bbox="862 436 1433 600"> <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</p> <p>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	<p>Response Time</p> <p>Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms).</p> <p>This commodity is populated via APM or DIF integrations.</p> <p>Transactions</p> <p>Transaction is a value that represents the per-second utilization of the transactions allocated to a given entity.</p> <p>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.</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).</p> <p>VMem Request</p> <p>If applicable, the virtual memory utilized by the container against the memory request.</p> <p>VCPU</p> <p>The virtual CPU (in mCores) utilized by the container against the CPU limit (if no limit is set, then node capacity is used).</p> <p>VCPU Request</p> <p>If applicable, the virtual CPU (in mCores) utilized by the container against the CPU request.</p>

Entity	Commodity
	<p>VCPU Throttling</p> <p>The throttling of container vCPU that could impact response time, expressed as the percentage of throttling for all containers associated with a Container Spec. In the Capacity and Usage chart for containers, <i>used</i> and <i>utilization</i> values reflect the actual throttling percentage, while <i>capacity</i> value is always 100%.</p>
Container Pod	<p>VMem</p> <p>The virtual memory utilized by the pod against the node physical capacity.</p> <p>VCPU</p> <p>The virtual CPU (in mCores) utilized by the pod against the node physical capacity.</p> <p>VMem Request</p> <p>The virtual memory request allocated by the pod against the node allocatable capacity.</p> <p>VCPU Request</p> <p>The virtual CPU (in mCores) request allocated by the pod against the node allocatable capacity.</p> <p>VMem Request Quota</p> <p>If applicable, The amount of virtual memory request the pod has allocated against the namespace quota.</p> <p>VCPU Request Quota</p> <p>If applicable, The amount of virtual CPU request (in mCores) the pod has allocated against the namespace quota.</p> <p>VMem Limit Quota</p> <p>If applicable, The amount of virtual memory limit the pod has allocated against the namespace quota.</p> <p>VCPU Limit Quota</p> <p>If applicable, The amount of virtual CPU limit (in mCores) the pod has allocated against the namespace quota.</p>
Container Spec	<p>VMem</p> <p>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).</p> <p>VCPU</p> <p>The virtual CPU (in mCores) historically utilized by any containers run for this workload against the CPU limit (if no limit is set, then node capacity is used).</p> <p>VMem Request</p> <p>If applicable, the virtual memory historically utilized by any containers run for this workload against the memory request.</p> <p>VCPU Request</p> <p>If applicable, the virtual CPU (in mCores) historically utilized by any containers run for this workload against the CPU request.</p>
Workload Controller	<p>VMem Request Quota</p>

Entity	Commodity
	<p>If applicable, The amount of virtual memory request the pod has historically allocated for this workload against the namespace quota.</p> <p>VCPU Request Quota</p> <p>If applicable, The amount of virtual CPU request (in mCores) the pod has historically allocated for this workload against the namespace quota.</p> <p>VMem Limit Quota</p> <p>If applicable, The amount of virtual memory limit the pod has historically allocated for this workload against the namespace quota.</p> <p>VCPU Limit Quota</p> <p>If applicable, The amount of virtual CPU limit (in mCores) the pod has historically allocated for this workload against the namespace quota.</p>
Namespace	<p>VMem Request Quota</p> <p>The total amount of virtual memory request for all pods allocated to the namespace against the namespace quota.</p> <p>VCPU Request Quota</p> <p>The total amount of virtual CPU request (in mCores) for all pods allocated to the namespace against the namespace quota.</p> <p>VMem Limit Quota</p> <p>The total amount of virtual memory limit for all pods allocated to the namespace against the namespace quota.</p> <p>VCPU Limit Quota</p> <p>The total amount of virtual CPU limit (in mCores) for all pods allocated to the namespace against the namespace quota.</p>
Virtual Machine (Node)	<p>VMem</p> <p>The virtual memory utilized by the node against the memory allocated to the hosting virtual machine.</p> <p>VCPU</p> <p>The virtual CPU utilized by the node against the CPU allocated to the hosting virtual machine.</p> <p>VMem Request</p> <p>The total amount of virtual memory allocated to pods with memory request against the allocatable capacity of the node.</p> <p>VCPU Request</p> <p>The total amount of virtual CPU allocated to pods with CPU request against the allocatable capacity of the node.</p> <p>Number Consumers</p> <p>The total number of pods running on the node against the maximum number of pods allowed.</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.</p>



Fabric Targets

A fabric target is a service that unites compute, network and storage access into a cohesive system. When you connect Workload Optimization Manager to fabric targets, it monitors the performance and resource consumption of your fabric interconnects, IO modules, chassis, and physical machines to assure application performance and utilize resources as efficiently as possible.

Once connected, Workload Optimization Manager discovers the blade servers that host the VMs, the chassis and datastores that provide resources to the blade servers, the IO modules and fabric interconnects that provide network resources, and the virtual datastores that provide storage resources to the VMs.

As part of this process, Workload Optimization Manager will stitch information from the fabric target and connected hypervisor targets to provide more granular data and information related to the applications and VMs running on the hypervisor-stitched blade servers. Combined with other targets, this information will support a top-down, application-driven approach to managing your environment.

For example:

When Workload Optimization Manager discovers that blade servers housed in a particular chassis have been designated as vCenter hosts, the supply chain stitches the blade servers and chassis to the corresponding vCenter datacenter to establish their relationship. When you set the scope to that datacenter and view the Health chart, you will see the blade servers in the list of hosts. In addition, when the datacenter is included in a merge policy (a policy that merges datacenters for the purpose of VM placement), the VMs in the blade servers apply the policy, allowing them to move between datacenters as necessary.

When you add application server targets, your applications and their individual components and services are discovered, enabling a view of your infrastructure from an individual application service to the physical hardware. Adding public cloud targets also allow for workloads to potentially migrate from your UCS infrastructure to the cloud, based on cost or available resources.

Supply Chain

Fabric targets add IO Module, Fabric Interconnect, Domain, and Chassis entities to the supply chain. The Chassis entities host physical machines (blade servers) – The physical machines also consume network connection commodities from IO Modules. The Fabric Interconnect supplies connectivity to the overall network, and also hosts the UCS Manager for UCS Targets. The Domain serves as the bottom-level pool of network resource, supplying the Fabric Interconnect.

Monitored Resources

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

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. ▪ Latency Latency is the measurement of storage latency utilized by the entity.
Blade	<ul style="list-style-type: none"> ▪ Net The utilization of data through the Blade's network adapters ▪ Treated as a Physical Machine of the underlying Hypervisor (see below) CPU, Mem, etc.
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
I/O Module	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port
Switch	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port ▪ PortChannel Amalgamation of ports with a shared net throughput and utilization

Cisco UCS Manager

The Cisco Unified Computing System (UCS) Manager is a management solution that participates in server, fabric, and storage provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection.

Workload Optimization Manager supports Cisco UCS Manager 2.2+.

UCS integrates all of these resources in a scalable multi-chassis platform to converge administration onto a single point. Managing these various entities on a network fabric with Workload Optimization Manager enables automation at the hardware level, including automated provisioning of hosts.

Prerequisites

- A service account Workload Optimization Manager can use to connect to Cisco UCS Manager

Entity Mapping

Workload Optimization Manager Mapping	UCS
Host	Server / Blade / Rack Unit
Chassis	Chassis
Datacenter	Datacenter
IO Module	IO Module
Switch	Fabric Interconnect
Network	Network

Adding UCS Targets

To add a UCS target, select the **Fabric and Network** category and choose one of the UCS Fabric options to match the version of UCS you want to manage. Then provide the following information:

- **Address:** The IP address of the Cisco UCS Manager
This gives access to the Fabric Manager that resides on the interconnect.
Workload Optimization Manager connects to the Cisco UCS Manager via the HTTP protocol by default.
- **Username/Password:** The credentials of the account Workload Optimization Manager will use to connect to Cisco UCS Manager.
specify the IP address and credentials for Cisco UCS Manager. Workload Optimization Manager discovers the fabric interfaces associated with that manager.

NOTE:

When providing a username, if the account is managed in Active Directory you must include the domain in case-sensitive spelling. For example, `MyDomain\john` is not the same as `mydomain\john`. For local user accounts, just provide the username.

- **Port Number:** The port Workload Optimization Manager will use to connect to the target.
- **Secure Connection:** When checked, Workload Optimization Manager will connect using SSL.

Supply Chain

Fabric targets add IO Module, Fabric Interconnect (Switch), and Chassis entities to the supply chain. The Chassis entities host hosts – The hosts also consume network connection commodities from IO Modules. The Fabric Interconnect supplies connectivity to the overall network, and also hosts the UCS Manager. The Domain serves as the bottom-level pool of network resource, supplying the Fabric Interconnect.

Actions

Workload Optimization Manager recommends actions for the various entities of the UCS Fabric Network as follows:

Entity Type	Action
Physical Machines	<ul style="list-style-type: none"> ▪ Start Physical Machine ▪ Provision Physical Machine ▪ Suspend Physical Machine
Chassis	<ul style="list-style-type: none"> ▪ Provision New Chassis
Fabric Interconnect	<ul style="list-style-type: none"> ▪ Add Port to Port Channel ▪ Remove Port from Port Channel ▪ Add Port
DPod (if Network Flow target is present)	<ul style="list-style-type: none"> ▪ Provision new DPod

Monitored Resources

Workload Optimization Manager monitors the following commodities of the UCS target:

Entity Type	Commodity
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared memory among VMs running on the host. ESX-only

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ CPU Ready The utilization of the PM’s allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Chassis	<ul style="list-style-type: none"> ▪ Power Electricity being consumed by the Chassis ▪ Cooling The percentage of the acceptable temperature range that is utilized by this chassis. As the chassis temperature nears the high or low running temperature limits, this percentage increases.
I/O Module	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port
Switch	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port ▪ PortChannel Amalgamation of ports with a shared net throughput and utilization
DPod (if Network Flow target is present)	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the DPod's memory reserved or in use ▪ CPU The utilization of the DPod's CPU reserved or in use ▪ Storage The utilization of the storage attached to the DPod ▪ Flow The utilization of the network flow capacity utilized by the DPod. This is divided into Flow1 (Low Cost) and Flow2 (Medium Cost) utilization

HPE OneView

HPE OneView is a management solution that streamlines provisioning and lifecycle management across compute, storage, and fabric. Through a unified API, infrastructure can be configured, monitored, updated, and re-purposed.

HPE OneView integrates all of these resources in a scalable multi-enclosure platform to converge administration onto a single point. Managing these various entities on a network fabric with Workload Optimization Manager enables automation at the hardware level, including automated provisioning of hosts.

Prerequisites

- A service account Workload Optimization Manager can use to connect to HPE OneView.
- HPE OneView 2.0 and compatible hardware.
- The **Banner Page** option for the user account should be disabled in the HPE OneView user interface.
- You should disable **Require Acknowledgement** for the user account in the HPE OneView user interface.

Adding HPE OneView Targets

To add a HPE OneView as a target, select the **Fabric** category and choose the HPE OneView radio button. Then provide the following information:

- **Address:**

The IP address of the HPE OneView target

This gives access to the Fabric Manager that resides on the VM.

Workload Optimization Manager uses the HTTPS protocol by default. In order to force the HTTP protocol, the Address must be entered in one of two ways. For example, an IP of 8.8.8.8 must be entered as http://8.8.8.8 or by using a specific HTTP port, such as 8.8.8.8:80.

- **Username/Password:**

The credentials of the account Workload Optimization Manager will use to connect to the HPE OneView target.

specify the IP address and credentials for HPE OneView. Workload Optimization Manager discovers the fabric interfaces associated with that instance.

NOTE:

When providing a username, if the account is managed in Active Directory you must include the domain in case-sensitive spelling. For example, MyDomain@john is not the same as mydomain@john. For local user accounts, just provide the username.

Supply Chain

Fabric targets add IO Module, Fabric Interconnect (Switch), Domain, and Chassis entities to the supply chain. The Chassis entities host physical machines – The physical machines also consume network connection commodities from IO Modules. The Fabric Interconnect supplies connectivity to the overall network. The Domain serves as the bottom-level pool of network resource, supplying the Fabric Interconnect.

NOTE:

For HPE OneView targets, the "Fabric Interconnect" entity exists as a false "Switch", and only as a pass-through for network resources. Unlike other fabric targets, such as UCS, there is no physical hardware that serves this function.

Actions

Workload Optimization Manager recommends actions for the various entities of the HPE OneView Fabric Network 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

Entity Type	Action
	<ul style="list-style-type: none"> Suspend VM Provision VM
Physical Machines	<ul style="list-style-type: none"> Start Physical Machine Provision Physical Machine Suspend Physical Machine
Fabric Interconnect	<ul style="list-style-type: none"> Add Port to Port Channel Remove Port from Port Channel Add Port
DPod (if Network Flow target is present)	<ul style="list-style-type: none"> Provision new DPod

Monitored Resources

Workload Optimization Manager monitors the following commodities of the HPE OneView target:

Entity Type	Commodity
Virtual Machine	<ul style="list-style-type: none"> Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. Storage Access Storage Access is the measurement of IOPS utilized by the entity. Latency Latency is the measurement of storage latency utilized by the entity.
Host	<ul style="list-style-type: none"> Memory (Mem) The utilization of the PM's memory reserved or in use CPU The utilization of the PM's CPU reserved or in use IO The utilization of the PM's IO adapters Net The utilization of data through the PM's network adapters Swap The utilization of the PM's swap space Balloon The utilization of shared memory among VMs running on the host. ESX-only

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Storage	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the datastore's capacity ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="text-align: center;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
I/O Module	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port
Switch	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port ▪ PortChannel Amalgamation of ports with a shared net throughput and utilization
DPod (if Network Flow target is present)	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the DPod's memory reserved or in use ▪ CPU The utilization of the DPod's CPU reserved or in use ▪ Storage The utilization of the storage attached to the DPod ▪ Flow The utilization of the network flow capacity utilized by the DPod. This is divided into Flow1 (Low Cost) and Flow2 (Medium Cost) utilization



Guest OS Processes Targets

Using WMI (Windows Management Instrumentation) or SNMP, Workload Optimization Manager can discover OS processes that are running on your VMs. Analysis creates an Application Component entity for each process.

Workload Optimization Manager does not generate actions for these discovered Application Components. For many types of applications, you can a target dedicated to that application. You should use Guest OS Process targets for cases where there is no other target you can use for the application. For more information, see [Applications and Databases Targets \(on page 9\)](#).

WMI

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

Prerequisites

- Underlying VM host targets added to Workload Optimization Manager (for discovery)
- A WMI user account Workload Optimization Manager can use to connect to the WMI targets. This account can either be an administrator user, or a non-administrator belonging to certain local user groups. For a full list of requirements, see [Creating a WMI User Account \(on page 74\)](#)
- WMI enabled on target VMs. For assistance in enabling WMI, see [Enabling WMI \(on page 74\)](#).

NOTE:

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

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

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

Adding WMI Targets

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

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

- **Username**
The username Workload Optimization Manager will use to connect to the WinRM service on the Virtual Machine hosting the application. This should not include the Active Directory domain.
- **Password**
The password Workload Optimization Manager will use to connect to the WinRM service on the Virtual Machine hosting the application.
- **Scope**
Workload Optimization Manager will search for Windows applications on the VMs found within the set scope. This scope can be set to the entire environment, single or multiple clusters, or particular virtual machines. Select this option and choose the scope for application discovery.

NOTE:

A single scope (target) must not contain more than 500 virtual machines. Workload Optimization Manager recommends utilization of multiple WMI targets for environments exceeding this limit. Contact Workload Optimization Manager Support for assistance in using multiple WMI targets.

- **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 database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

Enabling Guest Load Metrics

NOTE:

For environments that use WMI or SNMP probes, you must configure the Workload Optimization Manager installation to enable guest load metrics. Without this configuration, the WMI and SNMP probes will not display discovered guest load data.

To enable guest load metrics, you must edit the CR file and restart the platform:

1. Open `/opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml` for editing.
2. Enable guest metrics in the `properties` section of the file:

```
properties:
  repository:
    showGuestLoad: true
```

3. Save your changes to the file.
4. Apply your changes.

Execute the command:

```
kubectl apply -f /opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml
```

5. To restart the component, delete the repository pod.

Execute the command (where **{POD_ID}** is the ID that is appended to the pod in the `get pods` listing):

```
kubectl delete pod -n cwom repository-{POD_ID}
```

Supported Actions

NOTE:

The specific actions that Workload Optimization Manager recommends can differ, depending on the processes that Workload Optimization Manager discovers.

For other application components, Workload Optimization Manager can recommend actions based on the resources it can discover for the application. For example, Node.js® applications report CPU usage, so Workload Optimization Manager can generate vCPU resize actions and display them in the user interface.

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

Entity Type	Action
Applications	Without separate targets to discover Guest OS Processes or Application Servers, Workload Optimization Manager does not generate actions on applications. Instead, it generates resize actions on the host VMs. For on-prem environments, if host utilization is high enough on the physical machine running the application VM, Workload Optimization Manager can also recommend provisioning a new host.
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 ▪ Suspend VM ▪ Provision VM

Monitored Resources

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

Entity Type	Commodity
WMI/SNMP Application	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity.

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Latency <p>Latency is the measurement of storage latency utilized by the entity.</p>

Creating a WMI User Account

The service account Workload Optimization Manager uses to connect to WMI should be an Active Directory domain account.

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

1. Add the account to each of the following local groups:
 - WinRMRemoteWMIUsers__ (or Remote Management Users)
 - Performance Monitor Users

NOTE:

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

2. Grant permissions to the account.

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

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

Enabling WMI

Enabling Discovery

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

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

NOTE:

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

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

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

Execute `winrm quickconfig`.

3. Enable the WinRM *Negotiate* authentication scheme.

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

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

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

- Server setting:
`winrm set winrm/config/service '{@AllowUnencrypted="true"}'`
- Client setting:
`winrm set winrm/config/client '{@AllowUnencrypted="true"}'`

5. Restart the Remote Registry service.

After restarting the Remote Registry service, Workload Optimization Manager will discover the WMI targets.

SSL Connections

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

Some example SSL-related commands:

- List certificates on the system:

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

- Find Thumbprint for the default certificate:

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

- Enable SSL listener:

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

SNMP

Workload Optimization Manager will discover application and operating system resources using SNMP (Simple Network Management Protocol) on Linux hosts within the specified scope of your environment.

Prerequisites

- Underlying VM host targets added to Workload Optimization Manager (for discovery)
- SNMP service enabled/configured on target VMs to allow access to the OIDs required. The credentials used must allow access to the following OIDs:

- 1.3.6.1.2.1.1.1 - `sysDescr`
- 1.3.6.1.2.1.25.4.2.1.2 - `hrSWRunName`
- 1.3.6.1.2.1.25.5.1.1 - `hrSWRunPerfEntry`
- 1.3.6.1.2.1.25.3.3.1.2 - `hrProcessorLoad`
- 1.3.6.1.4.1.2021.4.5 - `memTotalReal`
- 1.3.6.1.4.1.2021.4.6 - `memAvailReal`
- 1.3.6.1.4.1.2021.4.11 - `memTotalFree`
- 1.3.6.1.4.1.2021.4.14 - `memBuffer`
- 1.3.6.1.4.1.2021.4.15 - `memCached`

NOTE:

For environments that include SNMP targets, Workload Optimization Manager can discover incorrect memory values for Linux systems. This can occur for the systems that use the net-snmp package, version 5.7.2-43.el7. You should use versions less than 5.7.2-43.el7, or greater than or equal to 5.7.2-47.el7.

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

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

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

Adding SNMP Targets

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

- Target Name

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

- Scope

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

NOTE:

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

- Community

The SNMP v2c community string Workload Optimization Manager will use to connect to the SNMP agent on the VM

- Port number

The Port number Workload Optimization Manager will use to connect to the SNMP agent on the VM

- Full Validation

When selected, Workload Optimization Manager will require all database servers hosted on the VMs in the selected scope to be a valid target. If Workload Optimization Manager is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

- Enable SNMPv3

When checked, Workload Optimization Manager will use SNMPv3 to connect to the virtual machines in the selected scope

- SNMPv3 Username/SecurityName

The Username/SecurityName that Workload Optimization Manager will use to connect to the virtual machines hosting the application

- SNMPv3 Enable Privacy

When checked, Workload Optimization Manager will encrypt using the privacy password to keep the connection private.

- SNMPv3 Authentication Password

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

- SNMPv3 Privacy Password

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

- SHA-1 For Authentication

When checked, Workload Optimization Manager uses SHA-1 Authentication to connect to the virtual machines in the selected scope. When unchecked, Workload Optimization Manager uses MD5

- AES For Privacy

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

NOTE:

Workload Optimization Manager does not support AES 192 or 256 for SNMP targets.

Enabling Guest Load Metrics

NOTE:

For environments that use WMI or SNMP probes, you must configure the Workload Optimization Manager installation to enable guest load metrics. Without this configuration, the WMI and SNMP probes will not display discovered guest load data.

To enable guest load metrics, you must edit the CR file and restart the platform:

1. Open `/opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml` for editing.
2. Enable guest metrics in the `properties` section of the file:

```
properties:
  repository:
    showGuestLoad: true
```

3. Save your changes to the file.
4. Apply your changes.

Execute the command:

```
kubectl apply -f /opt/turbonomic/kubernetes/operator/deploy/crds/charts_v1alpha1_xl_cr.yaml
```

5. To restart the component, delete the repository pod.

Execute the command (where **{POD_ID}** is the ID that is appended to the pod in the `get pods` listing):

```
kubectl delete pod -n cwom repository-{POD_ID}
```

Supported Actions

NOTE:

The specific actions that Workload Optimization Manager recommends can differ, depending on the processes that Workload Optimization Manager discovers.

For other application components, Workload Optimization Manager can recommend actions based on the resources it can discover for the application. For example, Node.js® applications report CPU usage, so Workload Optimization Manager can generate vCPU resize actions and display them in the user interface.

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

Entity Type	Action
Applications	Without separate targets to discover Guest OS Processes or Application Servers, Workload Optimization Manager does not generate actions on applications. Instead, it generates resize actions on the host VMs. For on-prem environments, if host utilization is high enough on the physical machine running the application VM, Workload Optimization Manager can also recommend provisioning a new host.
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 ▪ Suspend VM ▪ Provision VM

Monitored Resources

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

Entity Type	Commodity
WMI/SNMP Application	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.
Virtual Machine	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. ▪ Latency Latency is the measurement of storage latency utilized by the entity.

Enabling SNMP

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

Enabling SNMP v2

1. Obtain and install the SNMP server package for your Linux distribution and version. This is commonly called `net-snmp` in most package managers. Also install the `net-snmp-utils` package if it is available
2. Configure the SNMP daemon by editing `/etc/snmp/snmpd.conf` or running `snmpconf -i` at the command line in order to verify that:
 - SNMP v2c is enabled
 - A read-only community name has been set. This community name will be used by Workload Optimization Manager to communicate with the SNMP agent on the VM

NOTE:

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

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

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

Long form, VACM:

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

Short form, older syntax:

```
rocommunity mycommunity1 default system
```

Verify Your SNMP Setup

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

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

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

Enabling SNMP v3

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

Most package managers call this package `net-snmp`. To verify your configuration in step 5 below, you should also install the `net-snmp-utils` package on the Workload Optimization Manager VM.
2. Stop the SNMP service.
3. Create the SNMPv3 user.

Execute the following command:

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

For example, the command:

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

results in output similar to:

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

```
createUser snmpv3user SHA "snmpv3authPass" AES snmpv3encPass
```

adding the following line to `/etc/snmp/snmpd.conf`:

```
rouser snmpv3user
```

4. Verify your SNMP setup.

To verify the configuration, you can execute the following command on the Workload Optimization Manager VM:

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



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 Virtual Memory is the measurement of memory utilized by the entity.Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity.Storage Access Storage Access is the measurement of IOPS utilized by the entity.Latency Latency is the measurement of storage latency utilized by the entity.
Blade	<ul style="list-style-type: none">Net The utilization of data through the Blade's network adapters

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Treated as a Physical Machine of the underlying Hypervisor (see below) CPU, Mem, etc.
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
I/O Module	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port
Switch	<ul style="list-style-type: none"> ▪ NetThroughput Rate of message delivery over a port ▪ PortChannel Amalgamation of ports with a shared net throughput and utilization
Storage	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the datastore's capacity ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="margin-left: 40px;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity.

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="586 279 1518 348">▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. <li data-bbox="586 359 1518 457">▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array <li data-bbox="586 468 1518 564">▪ Latency The utilization of latency, computed from the latency of each device in the disk array.

Cisco HyperFlex

Cisco HyperFlex provides a hyperconverged platform that combines the networking and compute power of UCS with the storage capabilities of the HyperFlex HX Data Platform.

With the additional and refined storage information provided by HyperFlex, Workload Optimization Manager narrows the Desired State and recommends actions using the joint compute and storage information, gaining valuable insight into the interconnected nature of your environment.

For Workload Optimization Manager to make the most informed decisions, you must also add the compute nodes and any associated hypervisors. HyperFlex environments typically include:

- **Converged (HX) Nodes**
A combination of the cluster's storage devices into a single multi-tiered, object-based datastore.
- **Compute Nodes**
Cisco B or C series servers that make up the compute resources of the cluster, and are typically managed by a hypervisor.
- **Controller VMs**
Each HyperFlex node includes a Controller VM that intercepts and handles all the I/O from associated virtual machines. Workload Optimization Manager will not recommend actions for these VMs.

Prerequisites

- A service account on the controller VM that Workload Optimization Manager can use to connect to the HyperFlex environment

Adding HyperFlex Targets

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

- **Address**
The name or IP address of the HyperFlex controller VM. If a port is appended to the IP address, it must be a secure (HTTPS) port.
- **Username/Password**
Credentials for a user account Workload Optimization Manager will use to connect to the controller VM.
To specify a root user for the target, give the name as `local/root`. Note that earlier versions of Workload Optimization Manager accepted a username of `root`. If an existing HyperFlex target fails to validate after you have upgraded Workload Optimization Manager, then check to make sure the username for a root account is properly set to `local/root`.

If you use vSphere SSO for authentication to the HyperFlex target, you must prepend your username with `vc-`. For example, if your username is `Administrator@vsphere.local`, you would enter it as `vc-Administrator@vsphere.local`.

- Port Number

Port Number Workload Optimization Manager will use to connect to the controller VM.

Supply Chain

HyperFlex targets add Disk Array entities to the supply chain, and receive more granular information from the compute resources in your environment.

Entity Comparison

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

HyperFlex Name	Workload Optimization Manager Entity
Volume	Storage
HX Cluster	Disk Array

Supported Actions

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

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

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

NOTE:

For this target, Workload Optimization Manager discovers the HX Cluster as a Disk Array. When you see a provision action on this entity, you should determine which of the following is most relevant, based on your environment:

- Add disks to converged nodes
- Add a new converged node
- Add a new HX Cluster

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

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="text-align: center;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.

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 \(on page 89\)](#).

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 \(on page 89\)](#).

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 \(on page 91\)](#).

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.	

Entity Type	Can Be Automated	Recommendations only
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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p>NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.

Entity Type	Commodity
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 allocated CPU ▪ 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.

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**.



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 \(on page 168\)](#).

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 ▪ Reconfigure Storage ▪ Reconfigure Virtual Machine ▪ Suspend VM ▪ Provision VM
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 Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity.

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Latency Latency is the measurement of storage latency utilized by the entity.
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Storage	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the datastore's capacity ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="margin-left: 40px;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Datacenter	<p style="margin-left: 40px;">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 ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO

Entity Type	Commodity
	<ul style="list-style-type: none"> <li data-bbox="625 275 1052 304">The utilization of the PM's IO adapters <li data-bbox="586 317 670 346">▪ Net <ul style="list-style-type: none"> <li data-bbox="625 359 1263 388">The utilization of data through the PM's network adapters <li data-bbox="586 401 691 430">▪ Swap <ul style="list-style-type: none"> <li data-bbox="625 443 1052 472">The utilization of the PM's swap space <li data-bbox="586 485 711 514">▪ Balloon <ul style="list-style-type: none"> <li data-bbox="625 527 1498 556">The utilization of shared of memory among VMs running on the host. ESX-only <li data-bbox="586 569 756 598">▪ CPU Ready <ul style="list-style-type: none"> <li data-bbox="625 611 1490 661">The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Provider Virtual Datacenter	<ul style="list-style-type: none"> <li data-bbox="586 716 797 745">▪ Memory (Mem) <ul style="list-style-type: none"> <li data-bbox="625 758 1305 787">The utilization of the Datacenter's memory reserved or in use <li data-bbox="586 800 678 829">▪ CPU <ul style="list-style-type: none"> <li data-bbox="625 842 1263 871">The utilization of the Datacenter's CPU reserved or in use <li data-bbox="586 884 716 913">▪ Storage <ul style="list-style-type: none"> <li data-bbox="625 926 1279 955">The utilization of the storage attached to the Provider vDC.
Consumer Virtual Datacenter	<ul style="list-style-type: none"> <li data-bbox="586 991 797 1020">▪ Memory (Mem) <ul style="list-style-type: none"> <li data-bbox="625 1033 1305 1062">The utilization of the Datacenter's memory reserved or in use <li data-bbox="586 1075 678 1104">▪ CPU <ul style="list-style-type: none"> <li data-bbox="625 1117 1263 1146">The utilization of the Datacenter's CPU reserved or in use <li data-bbox="586 1159 716 1188">▪ Storage <ul style="list-style-type: none"> <li data-bbox="625 1201 1300 1230">The utilization of the storage attached to the Consumer vDC.

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 ManagerTargets \(on page 119\)](#).

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 \(on page 206\)](#).

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 \(on page 97\)](#)
- Configure remote management on each Hyper-V server. Refer to [Enabling Windows Remote Management \(on page 206\)](#)
- 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 \(on page 207\)](#)).
- 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 \(on page 86\)](#).

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, Re-size 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"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. ▪ Latency Latency is the measurement of storage latency utilized by the entity.
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space
Storage	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the datastore's capacity

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="text-align: center;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
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 ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space

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 infrastructure that target manages, and links it into a supply chain to deliver application performance management.

Prerequisites

- Target User Account

A user account that Workload Optimization Manager uses to connect to your vCenter Server and execute actions. For information about setting permissions for this account, see [Creating a User Account In vCenter Server \(on page 103\)](#).

General Considerations

Before you configure a vCenter Server target, you should consider the following:

- **Linked vCenters:**

For linked vCenters, you must add each vCenter Server separately so Workload Optimization Manager can communicate with each vCenter Server through a separate API endpoint.
- **Restricting Workload Optimization Manager Access to Specific Clusters:**

When you add a vCenter Server target, Workload Optimization Manager discovers all of the connected entities that are visible, based on the target account that it uses to connect to the vCenter Server target. If you have clusters or other entities you want to exclude from discovery, you can use the vSphere management client to the role of the Workload Optimization Manager account to `No access` for the given entities.
- **Shared Datastores:**

If you add more than one vCenter Server target that manages the same datastore, you can enable or disable datastore browsing to discover wasted files on the shared datastore:

- Enable datastore browsing:

To properly enable browsing, you must turn on the **Enable Datastore Browsing** option in the target configuration for each vCenter Server target that manages the shared datastore.

- Disable datastore browsing:

If you don't want datastore browsing over shared datastores, you must turn *off* the **Enable Datastore Browsing** option in the target configuration for each vCenter Server target that manages the shared datastore.

If set **Enable Datastore Browsing** differently for separate targets that manage the same datastore, datastore browsing can give inconsistent results for active and wasted files.

- **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:**

Enable Guest Metrics is enabled by default for all vCenter Server targets added to Workload Optimization Manager.

NOTE:

If your vCenter Server target was added prior to 8.7.3, this option is enabled as part of the update.

This target option instructs that Workload Optimization Manager request from the VC instance that it start collecting and reporting enhanced memory metrics from each virtual machine, using the VMware Tools 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 6.5 U3 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 \(on page 103\)](#).

- **VMware Cloud Deployments:**

VMware Cloud (VMC) is a managed cloud offering that provides dedicated VMware vSphere-based Software Defined Data Centers (SDDC). VMC on AWS hosts SDDC on AWS facilities. For on-premise VMC, SDDC runs on specific hardware that meets VMware specifications.

VMware Cloud on AWS, Azure VMware Services, Google Cloud VMware Engine, IBM Cloud for VMware and other cloud hosted VMware solutions are treated like any other VMware integration. The same familiar Workload Optimization Manager experience can be expected when using the VMware vCenter Server target integration against any vCenter Server target whether it is cloud hosted or not, provided you pay close attention to satisfying all of the vCenter target prerequisites, especially permissions and roles. For example, VMware Cloud bare metal host instances may be modeled in Workload Optimization Manager as Hyperconverged templates, assigned a cost, and used to plan the migrating from existing on-premises host clusters to cloud hosted VMware solutions. Once in the cloud, you can use Workload Optimization Manager to plan migrations from one cloud hosted VMware solution to another. Provision and Suspend Host actions generated by Workload Optimization Manager may be more readily actionable on cloud hosted VMware solutions and automated using custom Action Scripts if required. If you have any requirements or improvements that relate specifically to cloud hosted VMware solutions, submit an enhancement request through [IBM Ideas](#) for future consideration.

For VMC environments, you can use the Workload Optimization Manager vCenter Server target to manage workloads, but you should be aware of the following concerns:

- DRS Settings
 - VMC SDDC requires DRS to be set to migration level 3. You cannot change this setting.
- Datacenter Names

VMC SDDC datacenter names are always SDDC Datacenter. This is true even for multiple VMC targets. You can use the `Target Name` filter to search and group VMs across the set of VMC SDDC datacenters and clusters that share this name.

- Discovery
 - Workload Optimization Manager does not discover VMC service costs or subscription details.
 - The user interface does not show cost savings or investments.
- Migration Plans

Plans do not discover VMC bare metal hosts. You can use Workload Optimization Manager HCI Host Templates to represent the bare metal instances. Then the plan can determine the number of VMC instances you need to support your existing on-premise workloads.

Migrate to Cloud plans do not support VMC.

Adding vCenter Server Targets

To add vCenter Server 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 Server 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 Server 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 \(on page 105\)](#)).

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 Server vMotions in its recommendations.

To configure Workload Optimization Manager to support cross vCenter Server 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.

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, Move VM Storage, 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"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity.

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Latency Latency is the measurement of storage latency utilized by the entity.
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Storage	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the datastore's capacity ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="margin-left: 40px;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Datacenter	<p style="margin-left: 40px;">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 ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO

Entity Type	Commodity
	<ul style="list-style-type: none"> The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared of memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Provider Virtual Datacenter	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the Datacenter's memory reserved or in use ▪ CPU The utilization of the Datacenter's CPU reserved or in use ▪ Storage The utilization of the storage attached to the Provider vDC.
Consumer Virtual Datacenter	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the Datacenter's memory reserved or in use ▪ CPU The utilization of the Datacenter's CPU reserved or in use ▪ Storage The utilization of the storage attached to the Consumer vDC.

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:

vCenter Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ▪ Read-only role for all entity types Assign either Global permissions or permissions for the given vCenter Server instance to the target user or user group.

Workload Optimization Manager Functionality	Required Permissions
	<ul style="list-style-type: none"> ▪ Requirement to monitor VSAN and storage profiles In order to enable VSAN support and discover groups based on storage profiles, you must enable the Profile-driven storage view permission. This permission is <i>disabled</i> in the built-in readonly role.
Recommend Actions	<ul style="list-style-type: none"> ▪ Read-only role for all entity types Assign either Global permissions or permissions for the given vCenter Server instance to the target user or user group.
Wasted Storage Reporting	<ul style="list-style-type: none"> ▪ Datastore > Browse Datastore
Execute VM Move	<ul style="list-style-type: none"> ▪ Resources > Assign VM to Resource Pool ▪ Resources > Migrate Powered Off VMs ▪ Resources > Migrate Powered On VMs ▪ Resources > Modify Resource Pool ▪ Resources > Query Vmotion
Execute VM Storage Move	<ul style="list-style-type: none"> ▪ Datastore > Allocate Space ▪ Datastore > Browse Datastore ▪ Resources > Assign VM to Resource Pool ▪ Resources > Migrate ▪ Resources > Modify Resource Pool ▪ Resources > Move Resource Pool ▪ Resources > Query VMotion ▪ Virtual Machine > Change Configuration > Change resource ▪ Virtual Machine > Change Configuration > Change Swapfile placement
Execute VM Resize	<ul style="list-style-type: none"> ▪ Virtual Machine > Change Configuration > Change CPU count ▪ Virtual Machine > Change Configuration > Change Memory ▪ Virtual Machine > Change Configuration > Change resource ▪ Virtual Machine > Interaction > Reset ▪ Virtual Machine > Interaction > Power Off ▪ Virtual Machine > Interaction > Power On
Discover Tags	<ul style="list-style-type: none"> ▪ Global > Global tag You must also open ports 10443 and 7443 on the target server
Guest Metrics	<ul style="list-style-type: none"> ▪ Performance > Modify Intervals

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.



Orchestrator Targets

Workload Optimization Manager supports Action Script and ServiceNow orchestrator targets.

With orchestrator targets you can integrate Workload Optimization Manager actions with the orchestrator's application management process. For example, you can pass Workload Optimization Manager to a Change Request system for approval, and the system can pass the action back to Workload Optimization Manager for execution.

NOTE:

Workload Optimization Manager uses remote servers to execute ActionScript orchestration. Managing the processes remotely means that you do not install custom code on the Workload Optimization Manager server. This eliminates associated security risks at that point. However, you are responsible to maintain the security of your Action Script Server, and to ensure the integrity of your custom code.

ActionScript Server

Resource Requirements for the Server

The remote server can be a VM or a container. The capacity you configure for the server depends entirely on the processes you intend to run on it. Workload Optimization Manager does not impose any special resource requirements on the server.

Configuring Command Execution

To support execution of your scripts, you must install any software that is necessary to run the scripts. This includes libraries, language processors, or other processes that your scripts will invoke.

Workload Optimization Manager invokes the scripts as commands on the server. The server must run an SSH service that you have configured to support command execution and SFTP operations. At this time, Cisco has tested action scripts with the OpenSSH sshd daemon.

The standard port for SSH is 22. You can configure a different port, and provide that for admins who configure the server as an ActionScript target.

Note that an action script can invoke any process you have deployed on the remote server. You do not have to run scripts; however, you must be able to invoke the processes from the command line. The script manifest gives Workload Optimization Manager the details it needs to build the command line invocation of each script.

Configuring the ActionScript User Account

To execute the scripts on your server, Workload Optimization Manager logs on via a user account that is authorized to execute the scripts from the command line. You provide the user credentials when you configure the ActionScript target. To support this interaction, the user account must meet the following requirements:

- Public Key

The user must have a public key in the `.ssh/authorized_keys` file. When you configure the ActionScript target, you provide this as the Private Token for the target.

NOTE:

For the ActionScript target to validate, the permissions on the `authorized_keys` file must be set to 600.

- Security for the `.ssh` Directory

The ActionScript User should be the only user with authorized access. You should set file permissions to 700.

- Supported Shells

The ActionScript User shell can be either the Bourne shell (usually at `/bin/sh`) or the Bourne-Again shell (usually at `/bin/bash`). Workload Optimization Manager passes parameters as it invokes your scripts. At this time it only supports script execution through these shells.

Handling ActionScript Timeouts

Workload Optimization Manager limits script execution to 30 minutes. If a script exceeds this limit, Workload Optimization Manager sends a `SIGTERM` to terminate the execution of the process.

Note that Workload Optimization Manager does not make any other attempt to terminate a process. For example you could implement the script so it traps the `SIGTERM` and continues to run. The process should terminate at the soonest safe opportunity. However, if the process does not terminate, then you must implement some way to terminate it outside of Workload Optimization Manager. Note that a runaway process continues to use its execution thread. This can block other processes (action scripts or primary processes) if there are no more threads in the pool.

Obtaining Private Key String

In order to pass the private key as a string, newlines must be inserted appropriately with the `\n` character, as the JSON format will not preserve these required newlines. One method to insert these lines on *nix systems is the `awk` command. This command should follow the following format: `awk 'ORS="\n"' path/to/keyfile > path/to/modified/keyfile`. For example, `awk 'ORS="\n"' ~/.ssh/id_rsa > /tmp/privkey.txt`. You can then `cat` or `vi/vim` the file to copy the text and insert it into the API call.

Example private key file, before modification:

```
-----BEGIN RSA PRIVATE KEY-----
MIIEKQIBAAKCAgEAXeKUK6WLFmf8seI2Kztl4DoKKw0d6AG/39SH6or5SiNbAGkM
Umv8ABA0aCQqMF01j9V3kzGn71DIWWTxLOMBH06I01dvd9XYPVRJotGjpxus+xmV
GbeXxJsdHfXrbtMeLnpkwGC7gIu6W2o5Ez6YuLDFk+CefZhGMOtVLiB2Nwtzn33D
KBh96tIom9EqKseovtQAj5VYDwPeuTaWX1RGMTl1tsNGyzeEc4Z2sppT4VFVhsUV
L+voDZzrtMfFiHUIgJPH4pK9c/iPj0um7YzmyMyARftyCsUg7iUZH4N1NLpikXVy
xKiro6Quu7SZ4kcvr3ZF3970CON0hHzC3mFsvZvRFgYKzgc8/BcZ8ZzEUFfqbQy4
deJCU+jo86dYFL2Q8Rff2zsJenG0WbVRLtLGkc7RbMdBGjtt5roV5VAp+TyR9Hdd
Dc5FYvt4fAAMJtcmTIzdEAPsbVzDhkSajxRhOXvezJn+2Ddr34/9ZvFE7nJNgP/8
KOGJ5H1dgG5mc7H+Ly3x5dXzPhcVNSUdNGgKnevzFgMJaGByAg5AZc5HomiIwFKs
KSE/m8dSRoHrHmenGGmk1IMUnvCC1mKkbUEJ2V6R17x8zmM4DYNRSTwcSL
SZ7U/yDuwpHwzeozNeoekGKeC23+jgyzr9/mzgmPF08q5F+wVg/7n1DRSXsCAwEA
AQKCAgBPHMPbt25AM6gCEfcoF3/BOFMCXS8cpz5npRdl1IQHsNTb+CN1/Df2aoQZ
MoQva8RTKbuxfEoweLJwktuc33B+bK5lF+q3lSxrdvwtfPSZHYFqXejnXoupXjh
loS2fO7vE/2dr86mmLY7f98ro3X5OKzsrQdmiMdfHwaj11TLdJdgJy5Tpf8K9lRG
3zuleMcs66MBxGo9LoXntDwwBog9e8Rk+fMeHwt46EQiKXLPfNPdmoNa8kyGzLqr
PsbpRCjPrRX3rkTRzYjPeCZXiGJn2uJZmCpCxgbhTrGZmLVofC+WCCefeg2M0cHh
```

```
R1XSUjVtR7KtyvvzZcQah6p8r6SCASQFeVbb13FJVkBd0J8Hzk5e4N5Qk3bv9V4W
F7CC8dRV66kcAkf68ytLGXkQASgo5RPBDYSp8E2JDLF4m+9w6aouPlg//QxG5+XA
wqGJseTipiHbiAlcIGScpc5QCqab2SNddaDtIIpnp2Zmo5E10jutGeUAE3k+cnJ1
JIFCKRL0vsPk0xODSXBv2PMAUjMT65ue7uBkDqCtgHiSBXTk6DiWBjduDnoZlAmC
I8M4o42mSPs2uoKcngJ2fboY5iSsXLQ8VT5Tix8fk7VzBFmOvchBUSFlyx6o3/0l
0SV3Qrqz7tblXx+axQapIUEVbLV140k0jepvxF3HVELXRQd2eQKCAQEA8EzXro5j
N/VgDQJ9pGphwUyQ9mNBC1c1jviAHHskh8gwa4BxbLzZ+aUUoIx0/mAYK0rihN+
QJLI84c9tA62Au6Tq+CT60ad4SwxQ5m9mQZjq888T3dRssmY+f+5kYBzmaPHv18I
Sohh2tHqNh68T1KmlbCABahc0ppqJtjp0oi0kR7XUk5gEeDNNb65oDotuNytiVd4
cNEq/c08rqq2rTupLLjhVzP9lgVXL/SBw/HBjpmYJ3DMrcPXH84JeMshJay/OnW
eY3Y+hsghw37zRxeUtn5to0wVYayAmEhyXVing7EVV5wsyUaJQE25ZZWLjjjBHP
ZpLsSRc++uugnwKCAQEA0tBSQRdD0H+06Ydu9Tno1IUeilTQ8+VGx3esFlKtiDTW
cSEZSfZKbf/GQGHKbXLVXjCLhbFpBcVcCcojbmGsmTWRWxlySYPsw4tBiYYSk6Fo8
J2heYW3HQzyj23RjXIGhCUqY5wNWj1brU4CRxHtMyV1eYJWGMtaBfDm8VoBghu9g
FAyPkJs4WZUKUI7Sbg2vcJyLLP4XETPlfACmptLrPs8RCyUjJAJYxpSdaSjLC7r9
9gBusQ0lIGpClrx+/86NYhdKDWmnd0lyjXNmVwVpIpoSVblXTZK7xRBgcNTsd2eh
h1HLuhLXZ76SplvtNH6MQP46WoblounEx7ONnsddpQKCAQEA11VD339DQktLo+7/
0CHPFZHMjueSwaHEtMQzyLORrsbBYXGMTqh0km0aqM1V1mNpvVFYt2z5pHEIaibL
w4mHk0/pqVndtgwjkaA9R1aeQeqg3SKysd5jOtzCkX0mIIC5G2X5KHUtoQWSepHX
XPRnND1c2c3fMqysNndiklatDgIx5GrklZXwnKGth+2ZNipCHmw/IQxy8fWvae8n
oYzrW9Ugv5YhaUqT0jeDP9PP1s0QvrHN/13eQnM2szicPDC6XRdK7cKxRrfWJ3P
kobHfHh8iLuYblF0sGRW7EZk04MvNhmxxlYXoGOY0lFBNz1+45jEmOe0g7Fypz
021NswKCAQAHx/6xTXvs4PByeDhFPuz0IzG8Y2MPDZBuPmGQtV6IX4LmzBT140xg
YQd9Wjmrnf1YK6yKs/0o/fiAZ8LLPjPVEC4Pxxz03TJQjuaCsBniJaqaVd2ARYAK
wEtAB2y/dprZj/JS3JGuWsP24hd/UpnD+PlJ0VsVfqiK4EqWrClalJ+ivQln0UW
MWYUvyDPRRwLqQuODpblYKqAyrhmU/Sw7c2wPp+sbZW/Rq8Euae8+Br1oJB1YjAP
ITzx+WUqpgIb88Yc9qBiy8g5pmAdT0uTuGi3eN2Dotlu6QDpIHQQ/f3WbKoxLGS
t6EMGOqJfH5Ky9d1jfmD70j/IgaBLNstAoIBAQCkqJNuH7FSta/KQ5fJzFrD/Mil
Q6yfb9D92UEabBlEt22Zr7Ry5Wav4DRcYwV1HrBJQBow4Gc5WzCEnE1HvhkX519w
hccp1P+tyRf+h6ppYqdl/YaW2vEMTR1gvvtV8+2aMYTNRQJuTjJIsk1vWvX42IHq
xqKM9xqmKGEsW+6e7UheaXY8afkjf20/euGgzTuEPoy+IA8vKrSanLLWm4xdYjMV
i9PFbuSyRV6tJqQzXz58dS2ykL5YpTLM+ZFYzmaAqSJWULpsKaxxE2fnmniXOYVF
E4VxEPsgNYH1k2eqtGRQwYR0ML/fihHbJl1QnMr6o7/tK1jx+3PCZeRaVR4/
-----END RSA PRIVATE KEY-----
```

The same private key, after running the above command:

```
-----BEGIN RSA PRIVATE KEY-----\n
MIIJKQIBAAKCAgEAXeKUK6WLFMF8seI2KztL4DoKKw0d6AG/39SH6or5SiNbAGkM\n
Umv8ABA0aCQqMF0l j9V3kzGn71DIWWTxLOMBH06I01dvd9XYPVRJotGjpxus+xmV\n
GbeXxJsdHfXrbtMeLnpkwC7gIu6W2o5Ez6YuLDkf+CeFZhGMotVliB2Nwtzn3D\n
KBh96tIom9EqKseovtQAj5VYDwPouTaWx1RGMT1ltsNGyezEc4Z2sppT4VFVhsUV\n
L+voDZzrtMfFiHUIg jPH4pK9c/iPj0um7YzmyMyArftyCsUg7iUZH4N1NLpikXVY\n
xKiro6Quu7SZ4kcvr3ZF3970CON0hHzC3mFsvZvRFgYKzgc8/BcZ8ZzEUFfQbQY4\n
dejCU+jo86dYFL2Q8Rff2zJenG0WbVRLtLGkc7RbMdBGjtt5roV5VAp+TyR9Hdd\n
Dc5FYvt4fAAMJtcmTIzdeAPSbVzDhkSajxRhOXvezJn+2Ddr34/9ZvFE7nJNgP/8\n
KOGJ5H1dgG5mc7H+Ly3x5dXzPhcVNSUDNGgKnEzvFgMJAByAg5AZc5HomiIwFKs\n
KSE/m8dSRoHrHmenGGmk1IMUnvCC1mKkbUEJ2V6R17x8zmM4DYNRSTwCSL\n
SZ7U/yDuwpHwzeozNeoekGKeC23+jgyzr9/mzgmPF08q5F+wVg/7nlDRSXsCAWEA\n
AQKCAgBPHMPbt25AM6gCEfcoF3/BOFMCS8cpz5nprdlInIQHSNTb+CN1/Df2aoQZ\n
MoQva8RTKbuxfEoweLjwktuc33B+bK5lF+q3lSxrdvwtfPSZhYFqXeJnXOupXjh\n
loS2f07ve/2dr86mmLY7f98ro3X5OKzsrQdmiMdfHwajl1TLdJdgJy5TpF8K91RG\n
3zulEmcs66MbxGo9LoXntDwwBog9e8Rk+fMeHwt46EQiKXLPfNPdmoNa8kyGzLqr\n
PsbpRCjPrRX3rkTRzYjPeCZXiGJn2uJZmCpCxbghTrGZmLVofc+WcCeFeg2M0cHh\n
R1XSUjVtR7KtyvvzZcQah6p8r6SCASQFeVbb13FJVkBd0J8Hzk5e4N5Qk3bv9V4W\n
F7CC8dRV66kcAkf68ytLGXkQASgo5RPBDYSp8E2JDLF4m+9w6aouPlg//QxG5+XA\n
wqGJseTipiHbiAlcIGScpc5QCqab2SNddaDtIIpnp2Zmo5E10jutGeUAE3k+cnJ1\n
JIFCKRL0vsPk0xODSXBv2PMAUjMT65ue7uBkDqCtgHiSBXTk6DiWBjduDnoZlAmC\n
I8M4o42mSPs2uoKcngJ2fboY5iSsXLQ8VT5Tix8fk7VzBFmOvchBUSFlyx6o3/0l\n
0SV3Qrqz7tblXx+axQapIUEVbLV140k0jepvxF3HVELXRQd2eQKCAQEA8EzXro5j\n
N/VgDQJ9pGphwUyQ9mNBC1c1jviAHHskh8gwa4BxbLzZ+aUUoIx0/mAYK0rihN+\n
QJLI84c9tA62Au6Tq+CT60ad4SwxQ5m9mQZjq888T3dRssmY+f+5kYBzmaPHv18I\n
```

```
Sohh2tHqNh68T1KmLbCABahc0ppqJTjP00i0kR7XUk5gEeDNNb65oDotuNytiVd4\n
cNEq/cO8rqk2rTupLLjhVKzP9lgVXL/SBw/HBjPMyJ3DMrcPXH84JeMshJay/OnW\n
eY3Y+hsghw37zRxeuUtn5to0wVYayAmEhyXVing7Evv5wsyUaJQE25ZZWLjjiBHP\n
ZpLsSRc++uugnwKCAQEA0tBSQRdGOH+O6Ydu9Tno1IUeilTQ8+VGx3esFlKtiDTW\n
cSEZSfZKbf/GQgHKbxLVXjCLhbfPbCvCcojbmGSmtWRWxlySYPSW4tBiYJsk6Fo8\n
J2hEYW3HQzyj23RjXIghCUqY5wNWj1brU4CRxHtMyV1eYJWgmTAbfDm8VoBghu9g\n
FAyPkJs4WZUkUI7Sbg2vcJyLLP4XETPlfACMptLrPs8RCyUjJAjYxpSdaJLC7r9\n
9gBusQO1IGpClrx+/86NYhdKDWmnd0lyjXNvMVwPIpoSVblXTZK7xRBgcNTsd2eh\n
h1HLuhLXZ76SplvtNH6MQP46WoblounEx7ONnsddpQKCAQEAllVD339DQktLo+/ \n
0CHPFZHMjueSwaHEtMQzyLORrsbBYXGMTqh0kM0aqMlVlmNpvVFYT2z5pHEIaibL\n
w4mHkO/pgVnDtgwjkA9R1aEQeqg3SKysd5jOtzCXkQ0mIIC5G2X5KHUtOQWSepHX\n
XPRnND1c2c3fMqysNnDiklatDgIx5GrklZXwnKGth+2ZNipCHmw/IQxy8fWvae8n\n
oYZurW9Ugv5YhaUqTOjedP9PP1s0QvrHN/13eQnM2szicPDC6XRdK7cKxRrfWJ3P\n
kobHfHh8iLuYblFc0sGRW7EZk04MvNhmnxlYXoGOY01fBNzl+45jEmOe0g7Fypz\n
021NswKCAQAHx/6xTXvs4PByeDhFPuz0IzG8Y2MPDZBuPmGQtV6IX4LmzBT14Oxg\n
YQd9WjmRnflYK6yKs/0o/fiAZ8LLPjPVEC4Ppkz03TJQjuaCsBniJaqAvd2ARyAK\n
wEtAB2y/dprZj/JS3JGuWsP24hd/UpnD+PlJ0VsVfqiK4EquWrClalJ+ivQln0UW\n
MWYUvyDPRRwLqQuODpblYKgAyrhmU/Sw7c2wPp+sbZW/Rq8Euae8+BrloJB1YjAP\n
ITzx+WUqpgIb88Yc9qBiY8g5pmAdT0uTuGIi3eN2Dot1u6QDpihQQ/f3WbK0XLGs\n
t6EMGOqJfH5Ky9d1jfmD70j/IgaBLNstAoIBAQCkqJNuH7FSta/KQ5fJzFrD/Mil\n
Q6yfb9D92UEabBlEt2Zr7Ry5Wav4DRcYwVlHrBJQB0w4Gc5WZcEnE1HvhkX519w\n
hccp1P+tyRf+h6ppYqdl/YaW2vEMTR1gvvtV8+2aMYTNRQJuTjJIsk1vWvX42IhQ\n
xqKM9xqmKGEsw+6e7UheaXY8afkjf20/euGgzTuEPoy+IA8vKrsANLLWm4xDYjMV\n
i9PFbuSyRV6tJqQZxZ58dS2ykL5YpTLM+ZFYzmaQsJWULpsKAXxE2fnmniXOYVF\n
E4VxEPsgNYH1k2eqtGRQwYR0ML/fihHbJlLQnMr6o7/tKlJx+3PCZeRaVR4/\n
-----END RSA PRIVATE KEY-----\n
```

Adding ActionScript Targets to Workload Optimization Manager

NOTE: At this time, ActionScript targets must be added using the API. The ActionScript target appears in the UI as a selection, but should not be used. For more information about the Workload Optimization Manager API and how to use it to add targets, see "API Reference" in the *API Guide*.

The TargetApiInputDTO for this target has the following parameters:

nameOrAddress

IP or FQDN of the script execution server.

port

Port used to connect to the script execution server.

userid

Username used to connect to the script execution server.

privateKeyString

Path to the SSH private token corresponding to the user used to connect to the script execution server.

manifestPath

Path to the ActionScript manifest file on the script execution server.

Use the POST `https://10.10.10.10/api/v3/targets` request to add this target. Send the constructed TargetApiInputDTO as the body of the request.

Example Input:

```
{
  "category": "Orchestrator",
  "inputFields": [
    {
      "name": "nameOrAddress",
```

```

    "value": "10.10.10.10"
  },
  {
    "name": "port",
    "value": "22"
  },
  {
    "name": "userid",
    "value": "shepp"
  },
  {
    "name": "privateKeyString",
    "value": "-----BEGIN RSA PRIVATE
KEY-----\nMI IJKQIBAACKAgEAXeKUK6WLFMf8seI2KztL4DoKKw0d6AG/39SH6or5SiNbAGkM\nUmV8ABA0aCQqMF01j9V3kzGn71DIW
WTxLOMBH06I0ldvd9XYPVRJotGjpxus+xmV\nGbeXxJsdHfXrbtMeLnpkwGC7gIu6W2o5Ez6YuLDFKf+CefZhGMotVLIb2Nwtzn33D\nnKBh
96tIom9EqKseovtQAj5VYDwPeuTaWX1RGMt1ltsNGyecz4Z2sppT4VfVhsUV\nnL
+voDZrztMfFiHUIgJPH4pK9c/iPj0um7YzmyMyArfTyCsUg7iUZH4N1NLpikXVy\nnxKiro6Quu7SZ4kcvr3ZF3970CON0hHzC3mFsvZvRF
gYKzgc8/BcZ8ZzEUFfQbqY4\ndeJCU+jo86dYFL2Q8Rff2zsJenG0WbVRLtLGkc7RbMdBGjtt5roV5VAp+TyR9Hdd\nndc5FYvt4fAAMJtc
mTIzdEAPsbVzDhkSajxRhOXvezJn+2Ddr34/9ZvFE7nJNgP/8\nnKOGJ5H1dgG5mc7H+Ly3x5dXzPhcVNSUDNGKneVzFgMJaGByAg5AZc5
HOmiIwFKs\nnKSE/m8dSRoHrHmenGGm5TtA4lulIMUnvCC1mKkbUEJ2V6Rl7x8zmM4DYNRSTwcSL\nnSZ7U/yDuwpHweozNeoekGKeC23+j
gyzr9/mzgmPF08q5F+wVg/7nlDRSXSCAWEA\nnAQKCAgBPHMPbt25AM6gCefcoF3/BOFMCXS8cpz5npRdlnIQHSNTb+CN1/Df2aoQZ\nnMoQ
va8RTkbuxfEOweLjwktuc33B+bK5lF+q3lSxrdvwtfPSZHYFqXeJnXOupXjh\nnloS2f07vE/2dr86mmLY7f98ro3X5OKzsrQdmiMdfHwa
j1lTLDJdgJy5TpF8K9lRG\nn3zuleMcs66MBxGo9LoXntDwwBog9e8Rk+fMeHWt46EQiKXLpfNPdmoNa8kyGzLqr\nnPsbpRCjPrRX3rkTRz
YjPeCZXiGJn2uJZmCpCxbgTrGZmlVofC+WcCefeg2M0cHh\nnR1XSUjVtR7KtyvzZcQah6p8r6SCASQfEbb13fJvKbd0J8Hzk5e4N5Qk
3bv9V4W\nnF7CC8dRV66kAkf68ytLGXkQASgo5RPBDYSp8E2JDLF4m+9w6aouPlg//
QxG5+XA\nnwqGJsetiPiHbiALcIGsCpc5QCqab2SNddaDtIIPnp2Zmo5E10jutGeUAE3k+cnJl\nnJIFCkRL0vsPk0xODSxbV2PmaUjMT65u
e7uBkDqCtgHiSBXtK6DiWbJduDnoZlAmc\nnI8M4o42mSPs2uoKcngJ2fboY5iSsXLQ8VT5Tix8fk7VzBFmOvchBUSFlyx6o3/0l
\nn0SV3Qrqr7tb1Xx+axQapIUEVbLVl40kOjepvxf3HVELXRQd2eQKCAQEAE8EzXro5j\nnN/VgDQJ9pGphwUyQ9mNBC1c1jviAHHskh8gwa4
BxhBLzZ+aUuIx0/mAYK0rihN+\nnQJLI84c9tA62Au6Tq+CT60ad4SwxQ5m9mQZjq888T3dRssmY+f
+5kyBzmaPHv18I\nnSohh2tHqNh68T1KmLbCABahc0ppqJTjp0oi0kR7XUK5gEdNNb65oDotuNytivd4\nncNEq/c08rqq2rTupLLjhVKzP
9lgVXL/SBw/HBjPMyJ3DMrcPXH84JemshJay/OnW
\nneY3Y+hsghw37zRxeuUtn5to0wVYayAmEhyXVing7EVV5wsyUaJQE25ZZWLjjibHP\nnZpLsSRc+
+uugnwKCAQEAE0tBSQrDdOH+O6Ydu9Tno1IUEiltQ8+VGx3esFlktidTW\nncSEZSfZKbf/GQgHKbxLVXjclhbfPbCvCcojbmGsmtWRWxlyS
YPSW4tBiYJSk6Fo8\nnJ2heYW3HQzyj23RjXIghCUqY5wNWj1brU4CRxHtMyVleYJWgmTAbfDm8VoBghu9g\nnFAyPkJs4WZUkUI7Sbg2vcJ
yLLP4XETPlfACMptLrPs8RCyUjJAJYxpsDaSjLc7r9\nn9gBusQOlIGpClrx+/86NYhdKDWmnd0lyjXNvMVwPIpoSVblXTZK7xRBgcNTsd
2eh\nnhlHLuhLXZ76SplvtNH6MQP46WoblounEx7ONnsddpQKCAQEAl1VD339DQktLo+7/\nn0CHPFZHMjueSwaHEtMQzyLORrsbBYXGMTqh
0kM0aqMlV1mNpvVFYT2z5pHEIaibL\nnw4mHkO/pqVnDtqwjka9R1aEQeg3SKysd5jOtzCXkQ0mIIC5G2X5KHUtOQWSepHX\nnXPRnND1c2
c3fMqysNndiklatDgIx5GrklZXwnKGth+2ZnIpCHmw/IQxy8fWvae8n\nnoYzurW9Ugv5YhaUqTOjEDP9PP1s0QvrHN/13eQnM2szicPDC6
XRdK7cKxRrfWJ3P\nnokobHfHh8iLuYblFc0sGRW7EZk04MvNhmnnxlYXoGOY0lFBNz1+45jEmOe0g7Fypz\nn021NswKCAQAHx/6xTXvs4PB
yeDhFPuz0IzG8Y2MPDZBuPmGQtV6IX4LmzBT140xg\nnYQd9Wjmrnf1YK6yKs/0o/fiAZ8LLPjPVEC4Ppkz03TJQjuaCsBniJaqAvd2AR
yAK\nnWetAB2y/dprZj/JS3JGUwSP24hd/UpnD+PlJ0VsvfqiK4EquWrClalJ+ivQln0UW\nnMWYUvyDPRRwLqQuODpblYKgAyrhmU/Sw7c2
wPp+sbZw/Rq8Euae8+BrloJBlYjAP\nnITzx+WUqpgIb88Yc9qBiy8g5pmdt0uTuGi3eN2Dotlu6QDpHQ/f3WbKOXLGs\nt6EMGOqJf
H5Ky9d1jfmD70j/IgaBLNstAoIBAQCkqJNuH7FSta/KQ5fJzFrD/Mil
\nnQ6yfB9D92UEabBlEt22Zr7Ry5WAv4DRcYwVlHrBJQB0w4Gc5WzCEnE1HvhkX519w\nnhccp1P+tyRf+h6ppYqdl/YaW2vEMTR1gvtv8+
2aMYTNRQJuTjJisklvWvX42IhQ\nnxqKM9xqmKGEsW+6e7UheaXY8afkjf20/euGgzTuEPoy+IA8vKrSAnLLWm4xDYjMV\nni9PFbuSyRV6t
JqQzXz58dS2ykL5YpTLM+ZFYZmhAqSJWULpsKAxxE2fnmniXOYVF\nnE4VxEPsgNYH1k2eqtGRQWYR0ML/fihHbJl1QnMr6o7/tKlJx+3PC
ZeRaVR4/\nn-----END RSA PRIVATE KEY-----\n"
  },
  {
    "name": "manifestPath",
    "value": "/home/shepp/action-script/manifest.json"
  }
],
"type": "ActionScript"
}

```

The API will return the TargetApiDTO for the ActionScript target upon success, or an error message upon failure.

Example Response:

```

{
  "uuid": "73356329749216",

```

```

"displayName": "10.10.10.10-/home/shepp/action-script/manifest.json",
"category": "Orchestrator",
"inputFields": [
  {
    "displayName": "Name or Address",
    "name": "nameOrAddress",
    "value": "10.10.10.10",
    "isMandatory": true,
    "isSecret": false,
    "isTargetDisplayName": false,
    "valueType": "STRING",
    "description": "IP or FQDNS for the Script Execution Server",
    "verificationRegex": ".*"
  },
  {
    "displayName": "Port",
    "name": "port",
    "value": "22",
    "defaultValue": "22",
    "isMandatory": false,
    "isSecret": false,
    "isTargetDisplayName": false,
    "valueType": "STRING",
    "description": "Port to use for the Script Execution Server",
    "verificationRegex": ".*"
  },
  {
    "displayName": "User ID",
    "name": "userid",
    "value": "shepp",
    "isMandatory": true,
    "isSecret": false,
    "isTargetDisplayName": false,
    "valueType": "STRING",
    "description": "Userid to use to execute command on the Script Execution Server",
    "verificationRegex": ".*"
  },
  {
    "displayName": "Private Token",
    "name": "privateKeyString",
    "isMandatory": true,
    "isSecret": true,
    "isTargetDisplayName": false,
    "valueType": "STRING",
    "description": "SSH Private Token corresponding to the Userid",
    "verificationRegex": ".*"
  },
  {
    "displayName": "Script Path",
    "name": "manifestPath",
    "value": "/home/shepp/action-script/manifest.json",
    "isMandatory": true,
    "isSecret": false,
    "isTargetDisplayName": false,
    "valueType": "STRING",
    "description": "File Path to the ActionScript manifest file on the Execution Server",
    "verificationRegex": ".*"
  },
  {
    "displayName": "Public Host Key",
    "name": "hostKey",
    "isMandatory": false,
    "isSecret": false,
    "isTargetDisplayName": false,
    "valueType": "STRING",
  }
]

```

```

    "description": "Public key presented by the SSH server for host authentication; if not provided, the
    presented key will be accepted and integrated into the target definition for future operations",
    "verificationRegex": ".*"
  }
],
"lastValidated": "2020-07-28T20:46:22Z",
"status": "Validated",
"type": "ActionScript",
"readonly": false
}

```

Response headers

```
cache-control: no-cache, no-store, max-age=
```

Once the success response is received, the target is validated and ready for use.

Flexera One

Workload Optimization Manager integrates its management of workloads with Flexera One License Management. To do this, Workload Optimization Manager discovers the licenses and entitlements that you have configured in your Flexera environment. Workload Optimization Manager then creates groups and policies to represent these licenses and entitlements.

- For each license, Workload Optimization Manager creates a dynamic group.

These groups contain the entities that are affected by a Flexera license. For example, Workload Optimization Manager might create one group of all the VMs affected by a license, and another group of all the host machines affected by that license.

These groups are dynamic – If you allocate or remove entities (workloads, hosts, etc.) under a given license, Workload Optimization Manager adds or removes those entities in the given group. Even if you have a license with no entities allocated to it, Workload Optimization Manager creates an empty group for that license.

The name of each group begins with `Flexera`. The group name then includes the name of the license, plus the entity type. For example, you might see the groups, `Flexera: Microsoft SQL Server 2016 Standard VMs` and `Flexera: Microsoft SQL Server 2016 Standard Hosts`.

- For supported licenses, Workload Optimization Manager creates Placement and Automation policies.

A Flexera license can include placement constraints, or resource utilization constraints. Workload Optimization Manager can create policies to comply with those constraints:

- Placement Policies

When calculating workload placement, Workload Optimization Manager uses placement policies to comply with constraints such as keeping workloads on licensed hosts, or clearing unlicensed workloads off of licensed hosts when necessary.

For more information, see "Placement Policies" in the *User Guide*.

- Automation Policies

Among other settings, these policies can specify operational, utilization, and scaling constraints to match the constraints in your Flexera license policies.

For more information, see "Automation Policies" in the *User Guide*.

The name for each policy begins with the name `Flexera`. The policy name then includes the name of the license and constraint. For example, you might see a policy named `Flexera: Microsoft SQL Server 2016 Standard restrict to hosts`.

Workload Optimization Manager then applies these policies to the groups that correspond to the entity type and license name.

Workload Optimization Manager Policies for Flexera Licenses

Workload Optimization Manager generates policies to comply with the following types of Flexera licenses, where **Usage Rights** lists the configurations Workload Optimization Manager considers when creating policies:

License:	Usage Rights:	Notes:
Microsoft Server Processor	Datacenter Edition <pre>"coverInstallsOnVirtualMachines": true, "limitNumberOfVirtualInstalls": false, "limitVirtualInstallsIncludesHost": true, "numberOfAllowedVirtualInstalls": 1</pre>	Create Placement Policies that limit on-prem VM moves to the licensed host.
	Standard Edition <pre>"coverInstallsOnVirtualMachines": true, "limitNumberOfVirtualInstalls": true, "limitVirtualInstallsIncludesHost": false, "numberOfAllowedVirtualInstalls": 1</pre>	Create Placement Policies that disable VM move actions.
Microsoft Server/Management Core	Datacenter Edition <pre>"coverInstallsOnVirtualMachines": true, "limitNumberOfVirtualInstalls": false, "limitVirtualInstallsIncludesHost": false</pre>	Create Placement Policies that limit on-premise VM moves to the licensed host.
	Standard Edition <pre>"coverInstallsOnVirtualMachines": true, "limitNumberOfVirtualInstalls": true, "limitVirtualInstallsIncludesHost": false, "minimumNumberOfProcessors": 2</pre>	Create Placement Policies that disable VM move actions.
Microsoft Server Core	<pre>"coverInstallsOnVirtualMachines": false, "limitNumberOfVirtualInstalls": false, "limitVirtualInstallsIncludesHost": true, "minimumNumberOfThreadsPerVM": 4</pre>	<p>Workload Optimization Manager analysis assumes licensing is based on the host.</p> <p>Create Placement Policies that limit on-prem VM moves to the licensed host. Respect minimum vCPU thresholds.</p> <p>If there is a value for <code>minimumNumberOfThreadsPerVM</code>, create Automation Policies to enforce that minimum vCPU threshold for VMs.</p>

Flexera Permissions

The account that Workload Optimization Manager uses to connect with a Flexera target must be a Read-Only account with the following roles:

- View IT assets
- call APIs

Adding Flexera Targets

You can configure a target to a Flexera SaaS account. To add a Flexera target, select the **Orchestration > Flexera** option on the Target Configuration page and provide the following information:

- **Flexera Region**
The region for your Flexera account. Choose one of:
 - NAM
 - EU
- **Organization ID**
The unique identifier that defines your organization within Flexera One.
- **Refresh Token**
The credential that Workload Optimization Manager can use to generate Access Tokens as needed.
- **Proxy Host**
The address of the proxy used for this target. Only fill out proxy information if you connect to the Flexera 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.
- **Secure Proxy Connection**
When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Actions

Workload Optimization Manager does not generate unique actions for entities it discovers through a Flexera target. Instead, it uses the license specifications it discovers in Flexera to prevent actions that do not comply with license restrictions. It can also modify or generate actions to ensure license compliance.

When Workload Optimization Manager generates an action, the action details include a list of policies that constrain the action recommendation. In this list, policy names that begin with `Flexera:` indicate policies that implement the license constraints Workload Optimization Manager has discovered for Flexera. You can click to navigate to the policy definition. Note that you cannot see the policy definition, but you cannot edit it.

Move Virtual Machine Tomcat from [Host] to [Host]
PERFORMANCE

CPU Congestion

VIRTUAL MACHINE
Tomcat

CURRENT HOST - IMPACT FROM ALL ACTIONS
[Host]

STATE
Action can be accepted and executed immediately.

POLICIES

Name	Type
[HA]NodeRole::master [Kubernetes-DC11KubeTarget]	Placement
Flexera: Microsoft Windows Server 2012 Datacenter restrict to hosts	Placement

The policies that Workload Optimization Manager generates affect Workload Optimization Manager actions in the following ways:

Entity:	Policy Impact on Actions:
Virtual Machine	<ul style="list-style-type: none"> ▪ Resize vCPU Limit Resize actions to stay above a specified minimum. Workload Optimization Manager will not recommend a resize that falls below a specified minimum. However, if a VM is already below the specified minimum and shows no performance problems, it will not recommend a resize up. ▪ Move Workload Optimization Manager restricts licensed VMs to their matching licensed Hosts. If the host has capacity, it can move unlicensed VMs onto a licensed host. However, if the host sees bottlenecks, Workload Optimization Manager will move unlicensed VMs off of the host first. Move actions can also consolidate license consumption on hosts. Assume you have one host that supports two licenses (LA and LB), and another host that only supports one license (LA). Also assume a mix of workloads, some that require both LA and LB, and

Entity:	Policy Impact on Actions:
	some that only require LA. Move actions can consolidate the LA workloads on the LA host, leaving more resources for workloads that also require LB.
Host	<ul style="list-style-type: none"> ▪ Provision Host Workload Optimization Manager can recommend to provision a new licensed host. For example, assume your environment includes hosts that support two different licenses (LA and LB). Now assume you add more workloads that require only LB licenses. Workload Optimization Manager can recommend that you provision a new host that supports the LB license. ▪ Reconfigure Host Workload Optimization Manager can recommend Host Reconfigure actions to change the licensing on a host. Assume Workload Optimization Manager has executed actions to consolidate LB-only workloads on a host with LA and LB licenses. Workload Optimization Manager can recommend that you reconfigure that host to remove LA, and only support the LB license.

ServiceNow

You can configure Workload Optimization Manager policies that log Workload Optimization Manager actions in your ServiceNow instance, and that submit actions for approval in ServiceNow workflows.

For more information concerning orchestration, see "Action Orchestration" in the *User Guide*.

NOTE:

When creating the action orchestration policy as explained in the section above, the scope of the policy must match the scope of the ServiceNow target.

Prerequisites

- A ServiceNow user with the `web_service_admin` role and the custom role `x_turbo_turbonomic.user` that is created during installation that can communicate with Workload Optimization Manager via the REST API.

Adding ServiceNow Targets

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

- Address
Hostname of the ServiceNow instance without the `http` or `https` protocols. For example, `dev-env-266.service-now.com`.
- Username
Username for the account Workload Optimization Manager will use to connect to the ServiceNow instance
- Password
Password for the account Workload Optimization Manager will use to connect to the ServiceNow instance
- Client ID

- The Client ID Workload Optimization Manager will use if Use OAuth is checked
- Client Secret
 - The password Workload Optimization Manager will use if Use OAuth is checked
- Port
 - Port used to access the ServiceNow Instance
- Use OAuth
 - When checked, Workload Optimization Manager will use OAuth authentication to connect to the ServiceNow target
- Proxy Host
 - IP address of the proxy server
- Proxy Port
 - Port used to access the proxy
- Proxy User
 - Username for the account Workload Optimization Manager will use to connect to the proxy
- Port
 - Port used to access the ServiceNow Instance

ServiceNow Integration

In order to complete target addition, see the [Workload Optimization Manager Actions for ServiceNow](#) documentation.



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.

NOTE:

Different targets use different names to refer to Virtual Datacenters. In the Workload Optimization Manager supply chain, these entities are all represented by Consumer and Provider VDCs, as follows:

Workload Optimization Manager	vCenter Server	VMM
Consumer VDC	Resource Pool (Child)	Tenant or TenantQuota
Provider VDC	Resource Pool (Root)	Cloud

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 ▪ CPU The utilization of the Datacenter's CPU reserved or in use ▪ Storage The utilization of the storage attached to the Provider vDC.
Consumer Virtual Datacenter	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the Datacenter's memory reserved or in use ▪ CPU The utilization of the Datacenter's CPU reserved or in use ▪ Storage The utilization of the storage attached to the Consumer vDC.

Virtual Machine Manager

In a Virtual Machine Manager(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 \(on page 206\)](#)

- 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. For instructions on synchronizing the clock on the Workload Optimization Manager server, see "Synchronizing Time" in the *Installation Guide*.
- 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 \(on page 209\)](#).

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 \(on page 97\)](#).

- 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

Entity Type	Action
	<ul style="list-style-type: none"> ▪ Suspend VM ▪ Provision VM
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 Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Virtual Storage (VStorage) Virtual Storage is the measurement storage utilized by the entity. ▪ Storage Access Storage Access is the measurement of IOPS utilized by the entity. ▪ Latency Latency is the measurement of storage latency utilized by the entity.
Host	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the PM's memory reserved or in use ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only

Entity Type	Commodity
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 ▪ CPU The utilization of the PM's CPU reserved or in use ▪ IO The utilization of the PM's IO adapters ▪ Net The utilization of data through the PM's network adapters ▪ Swap The utilization of the PM's swap space ▪ Balloon The utilization of shared of memory among VMs running on the host. ESX-only ▪ CPU Ready The utilization of the PM's allocated ready queue capacity that is in use, for 1, 2, and 4 CPU ready queues. ESX-only
Provider Virtual Datacenter	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the Datacenter's memory reserved or in use ▪ CPU The utilization of the Datacenter's CPU reserved or in use ▪ Storage The utilization of the storage attached to the Provider vDC.
Consumer Virtual Datacenter	<ul style="list-style-type: none"> ▪ Memory (Mem) The utilization of the Datacenter's memory reserved or in use ▪ CPU The utilization of the Datacenter's CPU reserved or in use ▪ Storage The utilization of the storage attached to the Consumer vDC.



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 discounts, you enable Workload Optimization Manager to use that richer pricing information to calculate workload size and discount coverage for your cloud 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. 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.

AWS GovCloud Targets

When adding GovCloud targets, ensure the following prerequisites are met:

- The GovCloud account must be either a master account with organization access, or a member account belonging to a master account already added to Workload Optimization Manager.
- A billing target with organization access containing billing information for the GovCloud target must be added to Workload Optimization Manager in order to see RI information.
RI utilization information is identical to non-government targets as outlined above.
- Each account within the organization must have a unique name.

Optionally, you can add the AWS standard account that manages invoicing for a GovCloud account. When you add this standard account, it appears in the Top Accounts chart and shows the total projected billed cost for the GovCloud account and the standard account itself.

Supported Regions - AWS

Workload Optimization Manager supports discovery and management of entities in the following AWS regions:

Region Code	Region Name	Notes
af-south-1	Africa (Cape Town)	Requires Opt-In within AWS console
ap-south-1	Asia Pacific (Mumbai)	
ap-northeast-2	Asia Pacific (Seoul)	
ap-southeast-1	Asia Pacific (Singapore)	
ap-northeast-1	Asia Pacific (Tokyo)	
ap-southeast-2	Asia Pacific (Sydney)	
ap-east-1	Asia Pacific (Hong Kong)	Requires Opt-In within AWS console
ap-northeast-3	Asia Pacific (Osaka)	
ca-central-1	Canada (Central)	
eu-central-1	Europe (Frankfurt)	
eu-south-1	Europe (Milan)	Requires Opt-In within AWS console
eu-west-1	Europe (Ireland)	
eu-west-2	Europe (London)	
eu-west-3	Europe (Paris)	
eu-north-1	Europe (Stockholm)	
me-south-1	Middle East (Bahrain)	Requires Opt-In within AWS console
sa-east-1	South America (São Paulo)	
us-east-1	US East (N. Virginia)	
us-east-2	US East (Ohio)	
us-west-1	US West (N. California)	
us-west-2	US West (Oregon)	
us-gov-east-1	AWS GovCloud (US-East)	Restricted Access (GovCloud)

Region Code	Region Name	Notes
us-gov-west-1	AWS GovCloud (US-West)	Restricted Access (GovCloud)

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.
- GovCloud toggle
When enabled, this target will be added as an AWS GovCloud (US) target.
- 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:

NOTE:

Workload Optimization Manager does not support adding AWS targets via IAM roles for an instance that was deployed on-prem as an OVA or VHD image. If your instance was deployed in this way, you must add the target *without* using IAM Roles (see above).

Workload Optimization Manager deployed in Kubernetes (Red Hat OpenShift, Amazon EKS, GKE) supports adding AWS targets via IAM roles if the cluster configurations support an OIDC provider and webhook method. See [AWS Target IAM Role Requirements \(on page 211\)](#).

We recommend 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 about IAM Groups, see <http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>.

- 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 internet 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

If you run Workload Optimization Manager behind a proxy or firewall, then you must configure Workload Optimization Manager to allow unrestricted access to the following URLs. This enables the full set of features for AWS targets.

NOTE:

If your firewall performs IP based routing, be aware that Amazon reserves the right to change the DNS names on these endpoints over time. Also note, the AWS SDK is designed to use dynamic endpoints. It is possible that Amazon will change the endpoints the SDK depends on over time. As a result, the names of endpoints you must access might be different than the names listed here.

- Autoscaling
`autoscaling.{region-id}.amazonaws.com`
- AWS Bill
`{bucket-name}.s3.{region-containing-report}.amazonaws.com`
- CloudWatch
`monitoring.{region-id}.amazonaws.com`
- CloudWatch Events
`events.{region-id}.amazonaws.com`
- CloudWatch Logs
`logs.{region-id}.amazonaws.com`
- Cost Explorer
`ce.us-east-1.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`
- Performance Insights
`pi.{region-id}.amazonaws.com`
- Price List
`pricing.us-east-1.amazonaws.com`
- Relational Database Service (RDS)

`rds.{region-id}.amazonaws.com`

- **Resource Groups**
`resource-groups.{region-id}.amazonaws.com`
- **Savings Plans (Optional: Required if you use Savings Plans in your environment)**
`savingsplans.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. For more information, see:

- [Creating Cost and Usage Reports](#) in the AWS documentation
- [Setting up an Amazon S3 bucket for Cost and Usage Reports](#) in the AWS documentation

Enabling Collection of Memory Statistics

We highly recommend enabling collection of memory metrics in your environment. Enabling memory metrics allows Workload Optimization Manager to generate Scale actions to optimize VM memory usage. For Workload Optimization Manager to collect memory metrics, you must enable the collection of these metrics on the VMs in your environment.

For more information, see [Enabling Collection of Memory Statistics: AWS \(on page 198\)](#).

AWS Permissions

The account for an AWS target must have permissions that enable Workload Optimization Manager to monitor entities in your environment, recommend actions, and execute actions.

Generic Permissions

You can use generic AWS permissions to set up access for the account Workload Optimization Manager uses to access your target. This is an easy way to configure a target account, but you do not have full control over the access you grant to that account. For finer control, you should set the minimum permissions.

Permission Level	Required Permissions
Read-Only (monitoring and recommendations)	<ul style="list-style-type: none"> ▪ <code>AmazonEC2ReadOnlyAccess</code> ▪ <code>AmazonS3ReadOnlyAccess</code>

Permission Level	Required Permissions
	<ul style="list-style-type: none"> ▪ AmazonRDSReadOnlyAccess ▪ AWSConfigRoleForOrganizations (only required for consolidated billing with the master account)
Execute Actions	<ul style="list-style-type: none"> ▪ AmazonEC2FullAccess ▪ AmazonS3ReadOnlyAccess ▪ AmazonRDSFullAccess ▪ AWSConfigRoleForOrganizations (only required for consolidated billing with the master account)

Minimum Permissions

To explicitly control the access that you grant to Workload Optimization Manager, as a minimum the account for an AWS target must have the following permissions:

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ▪ autoscaling:DescribeAutoScalingGroups ▪ cloudwatch:GetMetricData ▪ cloudwatch:GetMetricStatistics ▪ cloudwatch:ListMetrics ▪ 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 ▪ elasticloadbalancing:DescribeTargetGroups ▪ elasticloadbalancing:DescribeTargetHealth ▪ elasticloadbalancing:DescribeLoadBalancers ▪ elasticloadbalancing:DescribeInstanceHealth ▪ iam:GetUser ▪ organizations:DescribeOrganization ▪ organizations:ListAccounts ▪ pi:GetResourceMetrics ▪ rds:DescribeDBInstances ▪ rds:DescribeDBClusters ▪ rds:DescribeDBParameters ▪ rds:ListTagsForResource

Workload Optimization Manager Functionality	Required Permissions
	<ul style="list-style-type: none"> ▪ rds:DescribeOrderableDBInstanceOptions ▪ servicecatalog:SearchProducts
Action Execution	<ul style="list-style-type: none"> ▪ autoscaling:SuspendProcesses ▪ autoscaling:ResumeProcesses ▪ autoscaling:DescribeLaunchConfigurations ▪ autoscaling>CreateLaunchConfiguration ▪ autoscaling>DeleteLaunchConfiguration ▪ autoscaling:UpdateAutoScalingGroup ▪ ec2:DescribeInstances ▪ ec2:DescribeVolumes ▪ ec2:DescribeVolumesModifications ▪ ec2:ModifyInstanceAttribute ▪ ec2:StopInstances ▪ ec2:StartInstances ▪ ec2:ModifyVolume ▪ ec2:DescribeInstanceStatus ▪ ec2:DescribeReservedInstancesOfferings ▪ ec2>DeleteVolume ▪ KMS:CreateGrant (required if your EC2 instances use encrypted ECB volumes) ▪ rds:ModifyDBInstance ▪ servicecatalog:DescribeProduct ▪ servicecatalog:ProvisionProduct ▪ servicecatalog:DescribeRecord ▪ servicecatalog:ListLaunchPaths
Savings Plans	<ul style="list-style-type: none"> ▪ savingsplans:DescribeSavingsPlans

To enable access to an AWS Billing target, the account must also include these permissions (see [AWS Billing Targets \(on page 134\)](#)):

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ▪ ce:GetReservationUtilization ▪ ce:GetSavingsPlansUtilizationDetails ▪ ce:GetSavingsPlansUtilization ▪ ce:GetSavingsPlansCoverage ▪ iam:GetUser ▪ organizations:DescribeOrganization ▪ organizations:ListAccounts ▪ s3:GetBucketAcl ▪ s3:GetObject ▪ sts:AssumeRole ▪ sts:AssumeRoleWithWebIdentity

Cloud Instance Family Support

In the user interface, you can see the instance types that Workload Optimization Manager supports.

1. Navigate to **Settings > Policies**.
2. In the Policy Management page, search for and click **Virtual Machine Defaults**.
3. In the Configure Virtual Machine Policy page:
 - a. Scroll down to the bottom of the page.
 - b. Click **Add Scaling Constraint**.
 - c. Choose **Cloud Instance Types**.
 - d. Click **Edit**.

The policy page shows supported tiers for each cloud provider. A tier is a family of instance types, such as *M1* for GCP, *a1* for AWS and *Basic_A1* for Azure. Expand a tier to see individual instance types and resource allocations.

Workload Optimization Manager considers all supported instance types when making scaling decisions for cloud VMs. If you want your VMs to *only scale to* or *avoid* certain instance types, create policies for those VMs.

Actions

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

- **Virtual Machine**
 - **Scale**
Change the VM instance to use a different instance type or tier to optimize performance and costs.
 - **Stop and Start** (also known as 'parking' actions)
Stop a VM for a given period of time to reduce your cloud expenses, and then start it at a later time.
 - **Discount-related actions**
If you have a high percentage of on-demand VMs, you can reduce your monthly costs by increasing discount coverage. To increase coverage, you scale VMs to instance types that have existing capacity. If you need more capacity, then Workload Optimization Manager will recommend actions to purchase additional discounts.

For details, see "Cloud VM Actions" in the *User Guide* and

BROKEN LINK!!! Parking: Stop or Start Cloud Resources

- **Database Server**
 - Scale**
Scale compute and storage resources to optimize performance and costs.
For details, see "Cloud Database Server Actions" in the *User Guide*.
- **Volume**
 - **Scale**
Scale attached volumes to optimize performance and costs.
 - **Delete**
Delete unattached volumes as a cost-saving measure.

For details, see "Cloud Volume Actions" in the *User Guide*.

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 Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Storage Amount The utilization of the datastore's capacity ▪ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ▪ Net Throughput Rate of message delivery over a port ▪ Net Throughput Inbound Rate of message received over a port ▪ Net Throughput Outbound Rate of message sent over a port ▪ I/O Throughput The throughput to the underlying storage for the entity ▪ Latency The utilization of latency allocated for the VStorage on the VM
Database Server	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Storage Amount Storage Amount is the amount of Amazon EBS storage utilized by the entity. ▪ Storage Access Storage Access is IOPS utilized by the entity. ▪ DB Cache Hit Rate (if available) DB cache hit rate is the measurement of Database Server accesses that result in cache hits, measured as a percentage of hits versus total attempts. A high cache hit rate indicates efficiency. ▪ Connections Connection is the measurement of Database Server connections utilized by applications.
Volume	<ul style="list-style-type: none"> ▪ Storage Access The percentage of the volume's capacity for storage access operations (IOPS) that is in use. ▪ IO Throughput

Entity Type	Commodity
	The percentage of the volume's capacity for IO throughput that is in use. <ul style="list-style-type: none"> IO Throughput Read The percentage of the volume's capacity for IO throughput Read that is in use. IO Throughput Write The percentage of the volume's capacity for IO throughput Write that is in use.

AWS Billing Families

Search... FILTER

13 Targets

- aa.aws.amazon.com AWS VALIDATED: JAN 28, 202... >
- ab.aws.amazon.com AWS VALIDATED: JAN 28, 202... >

A star symbol indicates a master account.

Expand to see details.

ab.aws.amazon.com VALIDATED: JAN 28, 202... ▾

ABC (010101010101) ← Master account

RELATED ACCOUNTS

Prod (111111000000) Test (000000111111)

TestABC (121212121212) ← Member accounts

Validated Target status Jan 28, 2021 5:06:41 PM Last Validated

A greyed name indicates a member account that you have not configured as a target.

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. Note that you can have one AWS Billing target per Workload Optimization Manager instance.

Cloud service providers can offer their own price lists, including special costs for services or discounts for workloads. However, Workload Optimization Manager does not discover these adjustments. For example, to reflect any discounted prices in the Workload Optimization Manager display and in Workload Optimization Manager analysis, you must manually configure those discounts. In Workload Optimization Manager, you configure such discounts via **Price Adjustments** for specific billing groups in your cloud environment.

Price Adjustments are needed for AWS to show proper discounted rates that customers have, even after adding the AWS Master Billing Account target. If Price Adjustments are not set, Workload Optimization Manager will show on-demand pricing, which results in incorrect cost numbers in actions and the UI. See "Price Adjustments" in the *User Guide* for more information.

NOTE:

Billing targets use Cost and Usage reports.

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. For more information, see:

- [Creating Cost and Usage Reports](#) in the AWS documentation
- [Setting up an Amazon S3 bucket for Cost and Usage Reports](#) in the AWS documentation

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 AWS 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 AWS 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"> ▪ ce:GetReservationUtilization ▪ ce:GetSavingsPlansUtilizationDetails ▪ ce:GetSavingsPlansUtilization ▪ ce:GetSavingsPlansCoverage ▪ iam:GetUser ▪ organizations:DescribeOrganization ▪ organizations: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.

Google Cloud Platform

Google Cloud Platform (GCP) provides a scalable infrastructure platform in the cloud. Workload Optimization Manager gains access to this infrastructure through a GCP [service account](#) with the appropriate permissions.

When you add a service account as a target, Workload Optimization Manager discovers the projects that define compute, storage, and networking resources for your workloads. It then creates a derived target for each discovered project. Derived targets are not directly modifiable within Workload Optimization Manager but can be validated like any other target.

Workload Optimization Manager discovers a broader resource hierarchy if you add a service account target with permissions to retrieve folders or your entire GCP organization.

Workload Optimization Manager uses cost data in its analysis to make accurate recommendations for your workloads. To enable cost discovery and monitoring, you must grant your service account "Billing Account Viewer" access to the related [billing accounts](#), and then add those billing accounts as targets.

NOTE:

When you add billing accounts, Workload Optimization Manager also discovers negotiated pricing and committed use discounts (CUD) for your workloads. The Discount Inventory chart in the user interface shows a list of discovered CUDs.

Supported Regions and Zones

Workload Optimization Manager supports discovery and management of workloads in all currently available GCP [regions and zones](#).

Permissions

To configure GCP targets, you create service accounts with roles that grant Workload Optimization Manager the permissions it needs to discover your GCP resources and costs, and to execute actions (optional). This section lists the minimum permissions you need.

NOTE:

For instructions on creating service accounts in a `gcloud shell` session, see [GCP Target Service Account \(on page 201\)](#) and [GCP Billing Target Service Account \(on page 205\)](#).

Workload Optimization Manager Functionality	Minimum Authorization Scopes	Minimum IAM Roles/Permissions
(Required) Resource discovery	<i>Project-level scopes and permissions</i>	
	One of the following scope sets: Set 1: <ul style="list-style-type: none"> ▪ https://www.googleapis.com/auth/cloudplatform-projects.read-only ▪ https://www.googleapis.com/auth/cloud-platform ▪ https://www.googleapis.com/auth/cloud-billing.readonly Set 2: <ul style="list-style-type: none"> ▪ https://www.googleapis.com/auth/cloudplatform-projects.read-only ▪ https://www.googleapis.com/auth/cloud-billing.readonly 	<ul style="list-style-type: none"> ▪ <code>resourcemanager.projects.get</code> ▪ <code>compute.regions.list</code> ▪ <code>compute.zones.list</code> ▪ <code>compute.machineTypes.list</code> ▪ <code>compute.machineTypes.get</code> ▪ <code>compute.disks.list</code> ▪ <code>compute.disks.get</code> ▪ <code>compute.diskTypes.list</code> ▪ <code>compute.instances.list</code> ▪ <code>compute.instances.get</code> ▪ <code>compute.instanceGroupManagers.list</code> ▪ <code>compute.instanceGroupManagers.get</code> ▪ <code>logging.views.list</code> ▪ <code>logging.views.get</code> ▪ <code>monitoring.services.get</code> ▪ <code>monitoring.services.list</code> ▪ <code>monitoring.timeSeries.list</code> ▪ <code>serviceusage.services.get</code>

Workload Optimization Manager Functionality	Minimum Authorization Scopes	Minimum IAM Roles/Permissions
	<p>.com/auth/compute</p> <ul style="list-style-type: none"> ▪ https://www.googleapis.com/auth/monitoring.read ▪ https://www.googleapis.com/cloud-billing.readonly 	
	<p><i>Organization-level scopes and permissions</i></p> <p>NOTE: To target specific folders, define a custom role at the organization level. It is not possible to define custom roles at the folder level.</p>	
	<p>One of the following scope sets:</p> <p>Set 1:</p> <ul style="list-style-type: none"> ▪ https://www.googleapis.com/auth/cloud-platform ▪ https://www.googleapis.com/auth/cloud-billing.readonly <p>Set 2:</p> <ul style="list-style-type: none"> ▪ https://www.googleapis.com/auth/cloud-platform.read-only ▪ https://www.googleapis.com/auth/iam.test ▪ https://www.googleapis.com/auth/cloud-billing.readonly <p>Set 3:</p> <ul style="list-style-type: none"> ▪ https://www.googleapis.com/auth/cloudplatfor- 	<ul style="list-style-type: none"> ▪ All project-level permissions ▪ resourcemanager.organizations.get ▪ resourcemanager.folders.get ▪ resourcemanager.folders.list ▪ resourcemanager.projects.list ▪ resourcemanager.projects.get ▪ billing.resourceAssociations.list

Workload Optimization Manager Functionality	Minimum Authorization Scopes	Minimum IAM Roles/Permissions
	<ul style="list-style-type: none"> morganization-s.readonly ▪ https://www.googleapis.com/auth/cloudplatform-folders.read-only ▪ https://www.googleapis.com/auth/cloudplatform-projects.read-only ▪ https://www.googleapis.com/auth/iam.test ▪ https://www.googleapis.com/auth/cloud-billing.readonly 	
<p>(Required) Cost discovery</p> <p>Workload Optimization Manager discovers billing families, billed costs, negotiated pricing, and committed use discounts.</p>	<ul style="list-style-type: none"> ▪ All project-level scopes ▪ https://www.googleapis.com/auth/cloud-billing.readonly ▪ One of the following scopes: <ul style="list-style-type: none"> – https://www.googleapis.com/auth/compute – https://www.googleapis.com/auth/cloud-platform 	<p>The service account should have the Billing Account Viewer role and the following permissions to the project that stores billing data:</p> <ul style="list-style-type: none"> ▪ All project-level permissions ▪ billing.resourceAssociations.list ▪ billing.accounts.list ▪ compute.commitments.list <p>For queries of billed costs and negotiated pricing via BigQuery, the service account should have the Billing Account Viewer role and the following permissions to the project that stores billing data:</p> <ul style="list-style-type: none"> ▪ bigquery.tables.get ▪ bigquery.tables.getData ▪ bigquery.tables.list ▪ bigquery.jobs.create
<p>(Optional) Action execution</p>	<ul style="list-style-type: none"> ▪ All project-level scopes ▪ One of the following scopes: <ul style="list-style-type: none"> – https://www.googleapis.com/auth/ 	<ul style="list-style-type: none"> ▪ All project-level permissions ▪ compute.instances.stop ▪ compute.instances.setMachineType ▪ compute.instances.start ▪ compute.disks.delete ▪ compute.zoneOperations.get

Workload Optimization Manager Functionality	Minimum Authorization Scopes	Minimum IAM Roles/Permissions
	<ul style="list-style-type: none"> compute – https://www.googleapis.com/auth/cloud-platform 	<ul style="list-style-type: none"> ▪ compute.regionOperations.get ▪ compute.globalOperations.get

Firewall Access

If you run Workload Optimization Manager behind a firewall, you must configure Workload Optimization Manager to allow unrestricted access to the following GCP URLs:

- https://cloudresourcemanager.googleapis.com
- https://cloudbilling.googleapis.com
- https://compute.googleapis.com
- https://monitoring.googleapis.com
- https://bigquery.googleapis.com

Enabling Collection of Memory Metrics

We highly recommend enabling collection of memory metrics in your GCP environment. With memory metrics, Workload Optimization Manager can generate actions that not only boost performance but also maximize your savings.

GCP collects these metrics via [Ops Agent](#). In order for Workload Optimization Manager to retrieve these metrics, you must install and configure Ops Agent on each VM that it monitors. See Ops Agent installation instructions [here](#), and configuration details [here](#).

NOTE:

GCP recommends using Ops Agent instead of its [legacy monitoring agent](#).

Enabling Required GCP APIs

For Workload Optimization Manager to discover your GCP environment and billing details, you must enable the following APIs:

- Cloud Resource Manager API
Creates, reads, and updates metadata for GCP resource containers.
- Compute Engine API
Creates GCP VMs and volumes.
- Cloud Billing API
Enables developers to manage billing for their GCP projects programmatically.
- BigQuery API
A data platform for customers to create, manage, share, and query data.

To enable these APIs:

1. Navigate the GCP Console to the library of APIs.
On the GCP Console home page, navigate to **APIs & Services > Library**.
2. Search for the API you want to enable.

In the API Library **Search** box, enter the name of the API you want to enable. Then press **Enter** to execute the search. Repeat these steps for each of:

- Cloud Resource Manager API
- Compute Engine API
- Cloud Billing API
- BigQuery API

3. Enable the given API.

In the list that appears, click the API name to navigate to that API page. If the API is not already enabled, click **Enable**. After you enable the given API, the console displays a details page for that API.

4. Navigate to the console Home page.

For each API you want to enable, navigate back to the home page and repeat these steps.

Adding Service Accounts as Targets

Once a service account has been properly configured for use with Workload Optimization Manager, you must add it as a target from the Target Configuration page.

Specify the following when adding the target:

- Display Name

The name that identifies the target in Workload Optimization Manager. This name is for display purposes only and does not need to match any name in GCP.

- Service Account Key (JSON)

The [service account key](#) for the account you want to manage. Paste the JSON object for the account key into this field.

- Proxy Information

Specify proxy information only if you connect to a GCP target via proxy.

- 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.

- Secure Proxy Connection

When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Adding Billing Accounts as Targets

To add a billing account target, specify the following:

- Target Name

The name that identifies the target in Workload Optimization Manager. This name is for display purposes only and does not need to match any name in GCP.

- Service Account Key

The [service account key](#) for the account associated with the billing account.

- GCP Project ID

The unique ID assigned to the project associated with the billing account. Costs accrued to this project are charged to the billing account you are adding.

- BigQuery Settings

BigQuery is a data warehouse that helps you manage GCP data. Workload Optimization Manager uses BigQuery resources to discover cost data for your environment. If you do not configure any of these fields, this target will not discover any cost data for Workload Optimization Manager analysis.

For more information, see:

- [BigQuery Resources](#)
- BigQuery Tables
 - [Example Queries for Cloud Billing Data Export](#)
 - [Schema - Standard Usage Cost Data](#)
 - [Schema - Pricing Data](#)

To configure your target to discover BigQuery data, specify values for the following fields.

- BigQuery Cost Export Data Set Name

The data set for billed costs. After you specify a data set, you must also specify the corresponding BigQuery Cost Export Table Name.

You can find the data set name in the GCP Billing dashboard under **Billing export / BIGQUERY EXPORT**.

- BigQuery Cost Export Table Name

This is the table of exported cost data. You can find the table name in the GCP BigQuery Explorer. Expand your project, and then expand the Cost Export Data Set Name.

- Enable Resource-Level Detail From Cost Export Table

When you configure Billing Export, you can enable **Detailed usage cost**. To expose this detailed data to Workload Optimization Manager, turn on this option and then give the name of the detailed data table in the **BigQuery Cost Export Table Name** field. You can find the table name in the GCP BigQuery Explorer.

NOTE:

Only turn on this option if you have enabled **Detailed usage cost**. If you want to provide a **Standard usage cost** table, do not turn on this option.

- BigQuery Pricing Export Table Name

This field automatically populates with the table name used in BigQuery, `cloud_pricing_export`. You need to provide a different name if you do any of the following:

- Use a different table for negotiated pricing
- Specify a value for the **BigQuery Pricing Export Data Set Name** field. In that case, you must also specify the corresponding pricing export table.

- BigQuery Pricing Export Data Set Name

The data set for pricing. You can find the data set name in the GCP Billing dashboard under **Billing export / BIGQUERY EXPORT**.

- Billing Account ID

The identifier of the Billing Account you want to target. This field is required if you configure **BigQuery Pricing Export Data Set Name** and **BigQuery Pricing Export Table Name**.

- Proxy Information

Specify proxy information only if you connect to a GCP target via proxy.

- 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.
- Secure Proxy Connection
When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Cloud Instance Family Support

In the user interface, you can see the instance types that Workload Optimization Manager supports.

1. Navigate to **Settings > Policies**.
2. In the Policy Management page, search for and click **Virtual Machine Defaults**.
3. In the Configure Virtual Machine Policy page:
 - a. Scroll down to the bottom of the page.
 - b. Click **Add Scaling Constraint**.
 - c. Choose **Cloud Instance Types**.
 - d. Click **Edit**.

The policy page shows supported tiers for each cloud provider. A tier is a family of instance types, such as *M1* for GCP, *a1* for AWS and *Basic_A1* for Azure. Expand a tier to see individual instance types and resource allocations.

Workload Optimization Manager considers all supported instance types when making scaling decisions for cloud VMs. If you want your VMs to *only scale to* or *avoid* certain instance types, create policies for those VMs.

Entity Mapping

After validating your GCP targets, Workload Optimization Manager updates the supply chain with the entities that it discovered. The following table describes the entity mapping between GCP and Workload Optimization Manager.

GCP	Workload Optimization Manager
Virtual Machine (VM) Instance	Virtual Machine (VM) NOTE: Workload Optimization Manager discovers GCP labels attached to VMs as tags. You can filter VMs by tags when you use Search or create groups. The Action Details page for a pending VM action also lists all the discovered tags.
Storage/Disk	Volume
Zone	Zone
Region	Region

NOTE:

GCP *projects*, *folders*, and *billing accounts* do not appear as entities in the supply chain. Use Search to scope to these resources. In Search, projects are grouped under Accounts, folders under Folders, and billing accounts under Billing Families.

For billing accounts, the latest billing data available from Google is always a few days old. As a result, billing-related charts in Workload Optimization Manager (such as the Billed Cost chart) do not have data for the current day. In addition, billing data shown in billing-related charts and GCP Billing reports might not match because Workload Optimization Manager uses UTC, while GCP uses local time. However, costs shown in both places are correct and reliable.

Monitored Resources

Workload Optimization Manager monitors the following resources for GCP workloads:

Entity Type	Commodity
Virtual Machine	<p>NOTE: Workload Optimization Manager does not monitor GCP machine types or CPU platforms that are currently in the preview/beta state.</p> <ul style="list-style-type: none"> ▪ Virtual Memory capacity utilized by the VM In order for Workload Optimization Manager to retrieve vMem metrics, you must install and configure Ops Agent on each VM that it monitors. ▪ Virtual CPU capacity utilized by the VM Workload Optimization Manager calculates capacity based on the normalized CPU frequency and the number of vCPUs for a given VM. Normalized CPU frequency takes into account performance variations seen in different models of a given CPU platform. Because frequency is normalized, charts might show utilization values that are slightly higher than 100% (for example, 100.03%) when capacity is fully utilized. ▪ Storage amount utilized by the VM ▪ IOPS and IO throughput (read/write) capacity utilized by the VM Workload Optimization Manager calculates capacity or uses GCP-published capacity data, depending on the VM's machine type and disk. <ul style="list-style-type: none"> – Shared-core machine types share a physical core and are used for running small, non-resource intensive apps. <ul style="list-style-type: none"> • For shared-core machine types with <i>standard disks</i>, Workload Optimization Manager uses capacity data that GCP publishes here. Note that there are no published values for e2-micro and e2-small, so Workload Optimization Manager assumes the e2-medium capacity for these machine types. • For shared-core machine types with <i>SSDs</i>, Workload Optimization Manager calculates capacity based on the observed maximum limit that can be achieved for IOPS and IO throughput, and uses the calculated capacity to analyze utilization more accurately. – For machine types that are <i>not</i> shared-core: <ul style="list-style-type: none"> • Workload Optimization Manager uses published capacity data and assumes that I/O block size is 16KB per I/O. • For machine types with persistent disks, Workload Optimization Manager assumes that the published capacity for the <i>SSD</i> disk type also applies to the <i>balanced</i> and <i>extreme</i> disk types.

Entity Type	Commodity
	When a VM is attached to at least one of these disk types, capacity is assumed to be the per-VM limit for the SSD disk type. When a VM is attached only to the <i>standard</i> disk type, capacity is the per-VM limit for the standard disk type. <ul style="list-style-type: none"> ▪ Net throughput (inbound and outbound) for a VM
Volume	Currently, Workload Optimization Manager does not monitor resources for GCP volumes. It only monitors their attachment state and then generates delete actions for unattached volumes.

Actions

Use the Potential Savings and Necessary Investments charts to view pending actions and evaluate their impact on your cloud expenditure.

- **Scale VM**

Scale VMs to optimize performance and costs. To generate accurate scaling actions, Workload Optimization Manager analyzes resource utilization percentiles and workload costs, and checks scaling constraints defined in policies.

Points to consider:

- Workload Optimization Manager can generate scaling actions for the following VMs, but cannot execute the actions automatically:
 - VMs with local SSDs

Workload Optimization Manager can recommend scaling to a machine type that supports local SSDs and the number of disks required by the VM, but blocks action execution due to prerequisite steps that you can only perform from GCP. You can view the prerequisite steps when you examine a pending action.
 - VMs configured with a [minimum CPU platform](#)

GCP instance type families can support multiple CPU generations. A specific VM may be configured with a minimum CPU platform to prevent it from scaling to instance types with incompatible CPUs. When you examine a pending action for such a VM, verify that the recommended instance type runs a compatible CPU. Once verified, manually execute the action from GCP.
- Workload Optimization Manager can recommend scaling VMs to instance types with existing CUD capacity. CUD purchase recommendations will be introduced in a future release.

NOTE:

Workload Optimization Manager does not recognize prioritized attributions you may have set for CUDs. For example, if you have prioritized all your CUD allotments for a single project, Workload Optimization Manager can still recommend actions to apply CUD to other projects in your environment.

- Since all GCP compute tiers have the same net throughput capacity, Workload Optimization Manager will not generate scaling actions in response to net throughput.
- Workload Optimization Manager does not recommend scaling actions for:
 - [Spot VMs](#)

NOTE:

Workload Optimization Manager discovers spot VMs, but does not recommend actions or monitor costs for these VMs.

- VMs running on [sole-tenant nodes](#)
- VMs with attached [GPUs](#)

- VMs in [managed instance groups](#)
 - VMs running [custom machine types](#)
- **Reconfigure VM**
GCP provides a specific set of machine types for each zone in a region. If you create a policy that restricts a VM to certain machine types and the zone it is currently on does not support all of those machine types, Workload Optimization Manager will recommend a reconfigure action as a way to notify you of the non-compliant VM. For example, assume Zone A does not support machine types for the M1 family. When a VM in that zone applies a policy that restricts it to M1, Workload Optimization Manager will recommend that you reconfigure the VM.
 - **Delete Volume**
Delete unattached GCP volumes as a cost-saving measure. Workload Optimization Manager generates an action immediately after discovering an unattached volume.

Workload Optimization Manager currently supports delete actions for [zonal](#) (single zone) persistent disks. The Potential Savings and Volume Summary charts show the savings you would realize if you execute these actions.

Points to consider:
 - Delete actions for [regional](#) persistent disks will be introduced in a future release. Currently, Workload Optimization Manager discovers these volumes, but does not show actions or costs in charts.
 - Workload Optimization Manager will never generate delete actions for local SSDs since they are always attached to VMs. GCP automatically deletes local SSDs when you delete the corresponding VMs.

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.

Azure Government Targets

Add the service principal account for Azure Government to enable the discovery and management of your Azure Government workloads. When you add the account, be sure to enable the **US Government** toggle to indicate that the account you are adding is for Azure Government.

In order for Workload Optimization Manager to discover custom pricing and reservations for Azure Government, you must add the corresponding Enterprise Agreement (EA) accounts. For details about EA accounts, see [Microsoft Enterprise Agreement \(on page 162\)](#).

NOTE:

Use [Azure Billing Targets \(on page 159\)](#) for non-government accounts and [Microsoft Enterprise Agreement \(on page 162\)](#) for government accounts.

To use Azure Government with a firewall or proxy, it must be configured to allow unrestricted access to certain URLs. For details, see [Firewall and Proxy Access \(on page 155\)](#).

Support for Azure App Service

Azure App Service is an HTTP-based service for hosting apps. With Azure App Service, app developers can easily create enterprise-ready apps and deploy them on a scalable and reliable cloud infrastructure.

Azure App Service offers several types of apps, including web apps, mobile apps, API apps, and logic apps. Each app runs as a set of *app instances* and is associated with a *plan* that defines compute resources (CPU, memory, and storage) available to the app.

When you add an Azure account:

- Workload Optimization Manager discovers all the plans in that account, except App Service Environment v3 I4, I5, and I6. Plans appear as 'Virtual Machine Specs' in the supply chain.
- For plans associated with *web apps*, Workload Optimization Manager discovers the related app instances. In the supply chain, app instances appear as 'App Component Specs'. Workload Optimization Manager generates actions to scale these plans to optimize app performance.
- For plans associated with the other types of apps, Workload Optimization Manager does not generate scale actions or discover the related app instances.
- For plans that are not associated with any type of app, Workload Optimization Manager generates delete actions as a cost-saving measure.

For details about scale and delete actions, see "Virtual Machine Spec" in the *User Guide*.

To discover plans and app instances, you must provide permissions to support all the actions you want to perform. For a list of permissions, see [Azure Service Principal and Subscription Permissions \(on page 150\)](#).

Supported Regions - Azure

Workload Optimization Manager supports discovery and management of entities in the following Azure regions:

Region Code	Region Name	Notes
eastus	East US	
eastus2	East US 2	
centralus	Central US	
northcentralus	North Central US	
southcentralus	South Central US	
westcentralus	West Central US	
westus	West US	
westus2	West US 2	
westus3	West US 3	
canadaeast	Canada East	
canadacentral	Canada Central	
brazilsouth	Brazil South	
brazilsoutheast	Brazil Southeast	
northeurope	North Europe	
westeurope	West Europe	
francecentral	France Central	
francesouth	France South	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)
ukwest	UK West	

Region Code	Region Name	Notes
uksouth	UK South	
germanynorth	Germany North	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)
germanywestcentral	Germany West Central	
norwayeast	Norway East	
norwaywest	Norway West	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)
switzerlandnorth	Switzerland North	
switzerlandwest	Switzerland West	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)
eastasia	East Asia	
southeastasia	Southeast Asia	
australiaeast	Australia East	
australiasoutheast	Australia Southeast	
australiacentral	Australia Central	
australiacentral2	Australia Central 2	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)
centralindia	Central India	
southindia	South India	
westindia	West India	
japaneast	Japan East	
japanwest	Japan West	
koreacentral	Korea Central	
koreasouth	Korea South	
uaecentral	UAE Central	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)
uaenorth	UAE North	
southafricanorth	South Africa North	
southafricawest	South Africa West	Access by request from Azure only (https://docs.microsoft.com/en-us/troubleshoot/azure/general/region-access-request-process)

Region Code	Region Name	Notes
		us/troubleshoot/azure/general/region-access-request-process
usgovarizona	USGov Arizona	Restricted Access (GovCloud)
usgovtexas	USGov Texas	Restricted Access (GovCloud)
usgovvirginia	USGov Virginia	Restricted Access (GovCloud)
usdodcentral	USDoD Central	Restricted Access (GovCloud) - Untested
usdodeast	USDoD East	Restricted Access (GovCloud) - Untested

Prerequisites

- 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.

- 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.

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](#).

Generic Azure Permissions

You can use generic Azure roles to set up access for the account Workload Optimization Manager uses to access your target. This is an easy way to configure a target account, but you do not have full control over the access you grant to that account. For finer control, you should set the minimum permissions.

Permission Level	Required Roles
Read-Only (monitoring and recommendations)	Reader plus Storage Account Contributor The use of a combined <code>Reader + Storage Account Contributor</code> role is the minimum privileged combination required for Workload Optimization Manager to discover and access metrics across your Azure environment. The <code>Storage Account Contributor</code> role is required to access the Storage Account keys and establish a connection in order to retrieve VM memory statistics.
Execute Actions	Owner or Contributor

Permission Level	Required Roles
	To enable actions, you can use either the <code>Owner</code> or <code>Contributor</code> role. <code>Contributor</code> is the least privileged role which enables Workload Optimization Manager to execute actions in your Azure environment.

Azure Service Principal and Subscription Permissions

For the Service Principal account, you must provide permissions to support all the actions you want to perform on any Subscription accounts.

Workload Optimization Manager Functionality	Required Permissions
Monitor workloads	<ul style="list-style-type: none"> ▪ <code>Microsoft.Capacity/reservationOrders/*/read</code> Monitor reservations pricing data. ▪ <code>Microsoft.Commerce/*/read</code> Monitor pricing data. ▪ <code>Microsoft.Commerce/RateCard/read</code> Discover pricing information from the Pay-as-you-go rate card. ▪ <code>Microsoft.Compute/*/read</code> Monitor status of compute resources to inform VM scaling decisions. ▪ <code>Microsoft.Consumption/pricesheets/read</code> Discover pricing information from an EA (Enterprise Agreement) Price Sheet. ▪ <code>Microsoft.ContainerService/managedClusters/read</code> Discover managed clusters. ▪ <code>Microsoft.ContainerService/managedClusters/agent-pools/read</code> Discover agent pools on managed clusters. ▪ <code>Microsoft.Insights/*/read</code> Monitor insights components. ▪ <code>Microsoft.Resources/subscriptions/read</code> The minimum for monitoring through a Service Principal target. The Azure-recommended <code>Reader</code> role and the <code>*/read</code> permission include this permission. ▪ <code>Microsoft.Network/networkInterfaces/read</code> Discover utilization on the network interface. ▪ <code>Microsoft.Network/publicIPAddresses/read</code> ▪ <code>Microsoft.Resources/subscriptions/*/read</code> ▪ <code>Microsoft.Sql/managedInstances/*/read</code> Discover managed instances. ▪ <code>Microsoft.Sql/servers/read</code> Discover SQL Server instances. ▪ <code>Microsoft.Sql/servers/databases/*/read</code> Discover metrics on SQL Server instances.

Workload Optimization Manager Functionality	Required Permissions
	<ul style="list-style-type: none"> ▪ (Optional) <code>Microsoft.Storage/storageAccounts/listkeys/action</code> Discover storage accounts as part of discovery for unmanaged disks.
Execute VM scaling, Start/Stop VM, Suspend VM	<ul style="list-style-type: none"> ▪ <code>Microsoft.Authorization/locks/read</code> ▪ <code>Microsoft.Compute/virtualMachines/write</code> ▪ <code>Microsoft.Compute/virtualMachineScaleSets/write</code> Enable actions per scale sets. ▪ <code>Microsoft.Network/networkInterfaces/join/action</code> VM can rejoin network after scaling. ▪ OPTIONAL: <code>Microsoft.KeyVault/vaults/deploy/action</code> For VM resize. Only needed if a VM is using Azure Key Vault, in order to deploy to the resized VM. ▪ <code>Microsoft.Compute/virtualMachines/deallocate/action</code> Stop VM to execute disruptive action. ▪ <code>Microsoft.Compute/virtualMachineScaleSets/deallocate/action</code> Stop VMs to execute disruptive action on scale set. ▪ <code>Microsoft.Compute/virtualMachines/start/action</code> Restart VM after stopping to execute disruptive action. ▪ <code>Microsoft.Compute/virtualMachineScaleSets/start/action</code> Restart VMs after stopping to execute disruptive action on scale set. ▪ <code>Microsoft.Compute/virtualMachines/powerOff/action</code> Suspend a VM.
Execute actions related to Microsoft Insights components	<ul style="list-style-type: none"> ▪ <code>Microsoft.Insights/AutoscaleSettings/Write</code> ▪ <code>Microsoft.OperationalInsights/workspaces/read</code> ▪ <code>Microsoft.OperationalInsights/workspaces/query/*/read</code> ▪ <code>Microsoft.OperationalInsights/workspaces/query/read</code>
Execute volume actions	<ul style="list-style-type: none"> ▪ <code>Microsoft.Compute/disks/write</code> Resize or change storage tier of disks. Reconnect disk to virtual machine after resize. ▪ <code>Microsoft.Compute/disks/delete</code> Delete unattached volume for managed disks. ▪ <code>Microsoft.Storage/storageAccounts/blobServices/containers/blobs/delete</code> Delete unattached volume for managed disks.
Scale SQL Database	<ul style="list-style-type: none"> ▪ <code>Microsoft.Sql/servers/databases/write</code> Scale databases per DTU or vCore pricing models. ▪ <code>Microsoft.Sql/servers/databases/pause/action</code>

Workload Optimization Manager Functionality	Required Permissions
	<p>Pause database.</p> <ul style="list-style-type: none"> Microsoft.Sql/servers/databases/resume/action <p>Restart database.</p>
Discover App Services (plans and app instances)	<ul style="list-style-type: none"> Microsoft.Relay/namespaces/HybridConnections/read Microsoft.Web/geoRegions/Read Microsoft.Web/serverfarms/Delete Microsoft.Web/serverfarms/Read Microsoft.Web/serverfarms/Write Microsoft.Web/serverfarms/metrics/read Microsoft.Web/serverfarms/sites/read Microsoft.Web/serverfarms/skus/read Microsoft.Web/serverfarms/usages/read Microsoft.Web/sites/metrics/read Microsoft.Web/sites/read Microsoft.Web/sites/slots/Read Microsoft.Web/sites/usages/read

For example, create a JSON file for your custom policy similar to the following, where `<Your_subscription_ID>` sets the role scope to the ID of the Azure subscription you want for this target:

```
{
  "Name": "Limited Access",
  "Description": "Limited access policy",
  "IsCustom": "true",
  "Actions": [
    "Microsoft.Authorization/locks/read",
    "Microsoft.Capacity/reservationOrders/*/read",
    "Microsoft.Commerce/*/read",
    "Microsoft.Compute/*/read",
    "Microsoft.Compute/disks/delete",
    "Microsoft.Compute/disks/write",
    "Microsoft.Compute/virtualMachines/deallocate/action",
    "Microsoft.Compute/virtualMachines/powerOff/action",
    "Microsoft.Compute/virtualMachines/start/action",
    "Microsoft.Compute/virtualMachines/write",
    "Microsoft.Compute/virtualMachineScaleSets/deallocate/action",
    "Microsoft.Compute/virtualMachineScaleSets/start/action",
    "Microsoft.Compute/virtualMachineScaleSets/write",
    "Microsoft.Consumption/*/read",
    "Microsoft.ContainerService/managedClusters/agentPools/read",
    "Microsoft.ContainerService/managedClusters/read",
    "Microsoft.Insights/*/read",
    "Microsoft.Insights/AutoscaleSettings/Write",
    "Microsoft.KeyVault/vaults/deploy/action",
    "Microsoft.Network/networkInterfaces/join/action",
    "Microsoft.Network/networkInterfaces/read",
    "Microsoft.Network/publicIPAddresses/read",
    "Microsoft.OperationalInsights/workspaces/query/*/read",
    "Microsoft.OperationalInsights/workspaces/query/read",
    "Microsoft.OperationalInsights/workspaces/read",
    "Microsoft.Relay/namespaces/HybridConnections/read",
```

```

"Microsoft.Resources/subscriptions/*/read",
"Microsoft.Resources/subscriptions/read",
"Microsoft.Sql/managedInstances/*/read",
"Microsoft.Sql/servers/databases/*/read",
"Microsoft.Sql/servers/databases/pause/action",
"Microsoft.Sql/servers/databases/resume/action",
"Microsoft.Sql/servers/databases/write",
"Microsoft.Sql/servers/read",
"Microsoft.Storage/storageAccounts/blobServices/containers/delete",
"Microsoft.Web/geoRegions/Read",
"Microsoft.Web/serverfarms/Delete",
"Microsoft.Web/serverfarms/metrics/read",
"Microsoft.Web/serverfarms/Read",
"Microsoft.Web/serverfarms/sites/read",
"Microsoft.Web/serverfarms/skus/read",
"Microsoft.Web/serverfarms/usages/read",
"Microsoft.Web/serverfarms/Write",
"Microsoft.Web/sites/metrics/read",
"Microsoft.Web/sites/read",
"Microsoft.Web/sites/slots/Read",
"Microsoft.Web/sites/usages/read"
],
"NotActions": [],
"AssignableScopes": [
  "/subscriptions/<Your_subscription_ID>"
]
}
    
```

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.

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](#).

Creating Client Secret Key and Permission

NOTE:

The administrator who creates the App registration must also create a Client secret key.

To create the secret key:

1. Navigate to the **Certificates and Secrets** section of your registered app.
2. Click + **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 since it **will not** be displayed again. 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.

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 Storage 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 (Service Principal or SPN) for this target must have permissions to manage the reservations. In most cases, `Reader` permissions are sufficient.

For [scoped reservations](#), Workload Optimization Manager supports the *subscription* and *shared* scopes. To discover these reservations, you must assign the Azure App registration `Reader` permissions to both the Reservation and the Reservation Order (order id). For details about reservation orders, see the Microsoft article, [Manage Reservations for Azure resources](#).

Workload Optimization Manager also discovers reservations scoped to a *resource group*, but treats them as shared (in the Discount Inventory chart, the scope for these reservations is shown as *Shared**). This could result in unreliable actions, such as scaling VMs within the resource group to other reservations, which could potentially increase costs. If you have existing reservations scoped to a resource group, be sure to change their scope in Azure before executing VM scale actions. For best results, change their scope to *shared*.

Adding Azure Targets

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

- Display Name
The display name you provide 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.
- US Government

Turn this on to configure the target as an Azure Government target.

- **Directory (Tenant) ID**
The ID of the tenant that contains subscriptions to be managed with Workload Optimization Manager.
- **Application (Client) 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

If you run Workload Optimization Manager behind a proxy or firewall, then you must configure Workload Optimization Manager to allow unrestricted access to the following URLs. This enables the full set of features for Azure Service Principal targets.

For Azure (Global):

- `ratecard.azure-api.net`
- `management.core.windows.net`
- `management.azure.com`
- `login.microsoftonline.com`
- `ratecard.blob.core.windows.net`
- `api.loganalytics.io`

For Azure Government:

- `ratecard.azure-api.net`
- `management.core.usgovcloudapi.net`
- `management.usgovcloudapi.net`
- `login.microsoftonline.us`
- `ratecard.blob.core.usgovcloudapi.net`
- `api.loganalytics.us`

Cloud Instance Family Support

In the user interface, you can see the instance types that Workload Optimization Manager supports.

1. Navigate to **Settings > Policies**.
2. In the Policy Management page, search for and click **Virtual Machine Defaults**.
3. In the Configure Virtual Machine Policy page:
 - a. Scroll down to the bottom of the page.
 - b. Click **Add Scaling Constraint**.
 - c. Choose **Cloud Instance Types**.
 - d. Click **Edit**.

The policy page shows supported tiers for each cloud provider. A tier is a family of instance types, such as *M1* for GCP, *a1* for AWS and *Basic_A1* for Azure. Expand a tier to see individual instance types and resource allocations.

Workload Optimization Manager considers all supported instance types when making scaling decisions for cloud VMs. If you want your VMs to *only scale to* or *avoid* certain instance types, create policies for those VMs.

Enabling Collection of Memory Statistics

We highly recommend enabling collection of memory metrics in your environment. Enabling memory metrics allows Workload Optimization Manager to generate Scale actions to optimize VM memory usage. For Workload Optimization Manager to collect memory metrics, you must enable the collection of these metrics on the VMs in your environment.

For more information, see [Memory Metrics Collection: Azure \(on page 200\)](#).

Actions

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

- **Virtual Machine**

- **Scale**

Scale up or down to template, based on VMem / VCPU. Change the VM instance to use a different instance type or tier to optimize performance and costs.

- **Move**

Execute intra-cloud moves of VMs.

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 must be a Linux VM with template configuration.

For details, see "Cloud VM Actions" in the *User Guide*.

- **Virtual Machine Spec (Azure App Service Plan)**

- **Scale**

Scale Azure App Service plans to optimize app performance or reduce costs, while complying with business policies.

- **Delete**

Delete empty Azure App Service plans as a cost-saving measure. A plan is considered empty if it is not hosting any running apps.

For details, see "Virtual Machine Spec Actions" in the *User Guide*.

- **App Component Spec (Azure App Service)**

None

Workload Optimization Manager does not recommend actions for App Component Specs, but it does recommend actions for the underlying Virtual Machine Specs. For details, see "Virtual Machine Spec Actions" in the *User Guide*.

- **Database**

Scale

- DTU Model

Scale DTU and storage resources to optimize performance and costs.

- vCore Model

Scale vCPU, vMem, IOPS, throughput and storage resources to optimize performance and costs.

For details, see "Cloud Database Actions" in the *User Guide*.

- **Volume**

- **Scale**

Scale attached volumes to optimize performance and costs.

– **Delete**

Delete unattached volumes as a cost-saving measure.

For details, see "Cloud Volume Actions" in the *User Guide*.

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) Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Storage Amount The utilization of the datastore's capacity ▪ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ▪ I/O Throughput The throughput to the underlying storage for the entity ▪ Latency The utilization of latency allocated for the VStorage on the VM
Virtual Machine Spec (Azure App Service Plan)	<ul style="list-style-type: none"> ▪ Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. ▪ Storage Amount Storage Amount is the measurement of Azure storage utilized by the entity. ▪ Number of Replicas Number of Replicas is the total number of VM instances underlying an App Service plan.
App Component Spec (Azure App Service)	<ul style="list-style-type: none"> ▪ Response Time Response Time is the elapsed time between a request and the response to that request. Response Time is typically measured in seconds (s) or milliseconds (ms). ▪ Virtual CPU Virtual CPU is the measurement of total CPU time utilized by a given app.
Database Server	<ul style="list-style-type: none"> ▪ Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity.

Entity Type	Commodity
Database	<p>NOTE: The resources that Workload Optimization Manager can monitor depend on the pricing model in place for the given database entity.</p> <ul style="list-style-type: none"> ▪ DTU Pricing Model <ul style="list-style-type: none"> – DTU DTU is the measurement of DTU capacity for the database. DTU represents CPU, memory, and IOPS/IO Throughput bundled as a single commodity. – Storage Storage is the storage capacity for the database. ▪ vCore Pricing Model <ul style="list-style-type: none"> – Virtual Memory Virtual Memory is the measurement of memory utilized by the entity. – Virtual CPU Virtual CPU is the measurement of CPU utilized by the entity. – Storage Access Storage Access is IOPS utilized by the entity. – Throughput Throughput is the utilization of transaction log write IO available to the entity. – Storage Storage is the storage capacity for the entity. <p>Workload Optimization Manager drives scaling actions based on the utilization of these resources, and treats the following limits as constraints when it makes scaling decisions:</p> <ul style="list-style-type: none"> ▪ Maximum concurrent sessions This is the maximum number of database connections at a time. ▪ Maximum concurrent workers This is the maximum number of database processes that can handle queries at a time.
Volume	<ul style="list-style-type: none"> ▪ Storage Access The percentage of the volume's capacity for storage access operations (IOPS) that is in use. ▪ IO Throughput The percentage of the volume's capacity for IO throughput that is in use. ▪ IO Throughput Read The percentage of the volume's capacity for IO throughput Read that is in use. ▪ IO Throughput Write The percentage of the volume's capacity for IO throughput Write that is in use.

Azure Billing

The Workload Optimization Manager Azure Billing target discovers your Azure billing account and *non-government* subscriptions. The target can access your billing data through your Microsoft Customer Agreement (MCA) Billing Account ID or Enterprise Agreement (EA) Enrollment Number.

IMPORTANT:

The Microsoft Enterprise Agreement target is nearing deprecation and will eventually be removed from the product.

If you have previously added a Microsoft Enterprise Agreement target to manage *non-government* subscriptions, be sure to move to the Azure Billing target *before* deprecation. To do this, remove the existing target from the user interface, and then add the Azure Billing target. For details about deprecation, see [Microsoft Enterprise Agreement \(on page 162\)](#).

If you have Azure *government* subscriptions, use the Microsoft Enterprise Agreement target until further notice. Currently, the Azure Billing target cannot be used to manage government subscriptions. For details about adding a Microsoft Enterprise Agreement target, see [Microsoft Enterprise Agreement \(on page 162\)](#).

Overview of the Azure Billing Target

The Azure Billing target discovers and supports the following:

- **Billing Organization**
The target discovers the billing account and related subscriptions associated with your MCA Billing Account ID or EA Enrollment Number.
- **Azure Reservations**
The target discovers all reservations that are charged under the billing account.
For MCA accounts, the target currently does not discover the percentage of workloads covered by reservations.
- **Billed Costs**
 - For EA accounts, the target reads data from a Cost Export that you set up in your environment. The data export is in CSV format, and contains all the cost and usage data that Azure Cost Management collects.
 - For MCA accounts, the target currently does not generate billed costs from a Cost Export.
- **Pricing**
The target discovers pricing for VMs, volumes, SQL databases (DTU and vCore), and App Service plans (web apps, logic apps, and function apps).
- **Actions**
The target uses pricing information when recommending the following actions:

Action	EA Accounts	MCA Accounts
Scale VM or volume	Yes	Yes
Optimize/Buy Azure reservations	Yes	No
Delete unattached volume	Yes	Yes
Scale SQL database (DTU and vCore)	Yes	Yes
Scale App Service plan (web apps only)	Yes	Yes
Delete unused App Service plan	Yes	Yes

Points to consider:

- Empty Azure EA subscriptions that are not incurring any charges will not stitch with the Azure Billing target, and a discrepancy will occur in the offer ID of the subscription. Once the subscription incurs a charge, the stitching occurs and the subscription should correctly associate with the Azure Billing target with the correct offer ID.
- Currently, when adding the Azure Billing target, the Top Billed Cost by Account, Top Billed Cost by Service, and Top Billed Cost by Service Provider charts do not display Azure billing data. This will be supported in a future release.
- For MCA accounts, Buy RI Actions are not generated and the Buy RI Plan is not supported. The Buy RI Plan will run but actions are not generated. Additionally, the Recommended Purchases widget is empty.

Prerequisites

Billing targets use Cost and Usage reports. In order for Workload Optimization Manager to display month-to-day spend, you must set up a daily Cost Export of actual cost on the Azure portal with an export type of month-to-date costs. The Cost Export must be created at the Billing Account scope. Subscription, Management Group, and Resource Group scopes are not supported.

We recommend creating a new Cost Export, even if you have an existing export that matches the setup noted here. Below is an example of a Cost Export setup.

- Export details
 - Metric
Actual cost (Usage and Purchases)
 - Export type
Daily export of month-to-date costs
- Storage
 - Use existing
 - Subscription
EA-Development
 - Storage account
turbocostexport
 - Container
cost-export-container
 - Directory
costExportDir

Azure Billing Target Permissions

To configure Azure Billing targets, you create service principal accounts with roles that grant Workload Optimization Manager the permissions it needs to discover your Azure billing resources and costs.

You must provide **Reader and Data Access** and **Storage Blob Data Reader** permissions for the storage account and container associated with the Cost Export.

Workload Optimization Manager Functionality	Required Permissions for Service Principal
Monitoring	<p>Enrollment Reader role (EA)</p> <ul style="list-style-type: none"> ▪ Only users with elevated access and the Enterprise Administrator role can apply the Enrollment Reader role to the service principal. ▪ The admin executing the put-assignment API must have elevated access in order to assign the Enrollment Reader at the tenant level. ▪ See Elevate access to manage all Azure subscriptions and management groups and Assign roles to Azure Enterprise Agreement service principal names in the Microsoft documentation. <p>Billing Reader role (MCA)</p> <ul style="list-style-type: none"> ▪ To assign the role, navigate to Cost Management + Billing > Access > Add role. Select the reader role and assign it to the service principal.

Firewall Access

If you run Workload Optimization Manager behind a firewall, you must configure Workload Optimization Manager to allow unrestricted access to the following URLs:

- `api.loganalytics.io`
- `login.microsoftonline.com`
- `management.azure.com`
- `[NAME_OF_THE_STORAGE_ACCOUNT_CONTAINING_THE_COST_EXPORT].blob.core.windows.net`

Adding an Azure Billing Target

NOTE:

Remove any corresponding non-government EA targets before adding an Azure Billing target.

For more information about adding an Azure Billing target, see [Azure Billing Target Setup](#).

To add an Azure Billing target, click **New Target** on the Target Configuration page, and select **Public Cloud > Azure Billing**.

Specify the following information:

- Billing Account ID

The ID of your billing account in Azure. You can find the billing ID in the **Cost Management + Billing** section of the Azure portal. The ID is different for MCA and EA accounts.

- MCA Billing Account ID

In the Azure portal, navigate to **Cost management + Billing** and then click your MCA billing account. The screen that displays includes a field for the Billing Account ID. For example, an ID could be:

```
095c24bb-6d32-5f7s-98e3-40123a4b56cd:2f98765e-a5b5
```

- EA Enrollment Number

In the Azure portal, navigate to **Cost management + Billing**, click to open the EA account, and navigate to **Properties**. The screen that displays includes a field for the Enrollment Number. For example, an Enrollment Number could be:

```
12345678
```

- **Cost Export Name**
The name of the cost export.
To find the cost export name in the Azure portal, navigate to **Cost management + Billing** and then select **Exports**. A list of the cost exports you have created displays. For more information about setting up cost exports, see [Prerequisites \(on page 160\)](#).
- **Directory (Tenant) ID**
The ID of the tenant that contains subscriptions to be managed with Workload Optimization Manager. This should match the Directory ID that you give for the associated Service Principal target.
- **Display Name**
The display name you provide to identify the target in the Target List. This name is for display purposes only and does not need to match any name in Azure.
- **Application (Client) ID**
The Client/App ID found in App registration. This ID gives Workload Optimization Manager access to resources in your Azure subscription, and should match the Client/App ID that you specify for the associated Service Principal target.
- **Client Secret Key**
The secret key for the App Registration. This should match the secret key ID that you give for the associated Service Principal target.
- **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.
- **Secure Proxy Connection**
When checked, Workload Optimization Manager will connect to the proxy via HTTPS.

Note that it takes 24 hours after successfully adding a target for any Azure reservation information to display in the user interface.

Actions

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

Monitored Resources

Workload Optimization Manager does not monitor resources for Azure Billing targets.

Microsoft Enterprise Agreement

IMPORTANT:

The Microsoft Enterprise Agreement target is nearing deprecation and will eventually be removed from the product.

Who is impacted?

Deprecation and removal impact customers who have previously added, or plan to add, a Microsoft Enterprise Agreement target to manage Azure *non-government* subscriptions.

- Before deprecation, customers who have previously added a Microsoft Enterprise Agreement target are advised to remove the target from the user interface, and then add the Azure Billing target.
- Customers planning to add a Microsoft Enterprise Agreement target should use the Azure Billing target instead.

The Azure Billing target leverages the new cost management APIs and cost exports recommended by [Microsoft](#). For instructions on adding the Azure Billing target, see [Azure Billing \(on page 159\)](#).

For impacted customers, the Microsoft Enterprise Agreement target will be deprecated in version **3.6.6**. Details about the removal will be available later.

When the Microsoft Enterprise Agreement target reaches deprecation:

- No new features will be introduced.
- The following charts may not show historical cost data initially. Customers who notice any data loss in their environment should note that the charts will rebuild the data over time.
 - Workload Cost Breakdown
 - Cost Breakdown by Tag
 - Top Billed Cost by Account
 - Top Billed Cost by Service
 - Billed Cost by Service Provider
- The target will continue to be available for customers to add or use, and will continue to discover pricing data.
- Customers can continue to report problems with the target. Workload Optimization Manager Support will answer questions, help with configuration issues, and provide any known workarounds.

Who is not impacted?

Deprecation and removal do *not* impact customers who have previously added, or plan to add, a Microsoft Enterprise Agreement target to manage Azure *government* subscriptions. Customers with these subscriptions should use the Microsoft Enterprise Agreement target until further notice. Currently, the Azure Billing target cannot be used to manage government subscriptions.

Overview of the Microsoft Enterprise Agreement Target

You can configure Workload Optimization Manager to manage *Azure government* subscriptions within the context of a Microsoft Enterprise Agreement (EA). An EA target enables Workload Optimization Manager to use custom pricing and discover reservations. When you configure an EA target, Workload Optimization Manager uses that richer pricing information to calculate workload size and reservation 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 principal target. For more information about service principal targets, see [Adding Azure Targets \(on page 146\)](#).

Prerequisites

- Microsoft Azure EA access key
- Your Microsoft Azure EA enrollment number
- Enable access to your Azure Billing Data

You can enable access to costs in the Azure portal or in the EA portal. For complete information, see [Assign access to Cost Management data](#) in the Microsoft documentation.

Azure Portal:

1. Log into to the Azure portal (<https://portal.azure.com>).
You should log in with an enterprise administrator account.
2. Open your Billing Account.
Navigate to **Cost Management + Billing | Billing scopes**. Select your billing account from the list of available accounts.
3. Configure cost access.
In the **Settings** group, select **Policies**. Then turn on **DEPARTMENT ADMINS CAN VIEW CHARGES** and **ACCOUNT OWNERS CAN VIEW CHARGES**.

EA Portal:

1. Log into to the Azure portal (<https://ea.azure.com>).
You should log in with an enterprise administrator account.
2. Select **Manage** in the left-hand navigation pane
3. Configure cost access.
For the cost management scopes that you want to provide access to, enable the charge option to **DA view charges** and/or **AO view charges**.

NOTE:

After you configure cost access, most scopes also require Azure role-based access control (Azure RBAC) permission configuration in the Azure portal.

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.

If you perform self-service exchanges for your reservations, Workload Optimization Manager does not discover the new charges for the exchanged reservations through the Azure EA target. To track the charges after you have exchanged reservations, ensure you have an Azure subscription target for the affected scope of Azure workloads, and that subscription has read access to reservations 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
The EA enrollment number (found in your EA admin account at ea.azure.com)
- API Key
The API Access Key for the EA (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 reservations 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 \(on page 146\)](#).

Azure Enterprise Agreements

The screenshot shows the configuration of an Azure Enterprise Agreement (EA) target in the Workload Optimization Manager. It is divided into two main sections: the EA target configuration and the Service Principal target configuration.

1) Enterprise Agreement (EA) target

- Target: **Azure-EA** (Azure EA)
- EA ENROLLMENT#: 11111111
- RELATED SUBSCRIPTIONS:
 - EA - Prod (x1x1x1x1-1a1a-x1x1-1a1a-x1x1x1x1x1)
 - EA - Test (x1x1x1x1-1a1a-x1x1-1a1a-x2x2x2x2x2)
 - x1x1x1x1-1a1a-x1x1-1a1a-x3x3x3x3x3

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.

- Target: **core.windows.net** (Azure Service Principal)
- CLOUD TYPE: Global
- TENANT ID: y1y1y1y1-1b1b-z1z1-1y1y-x1x1x1x1x1x1
- CLIENT ID: d1d1d1d1-1b1b-z1z1-1y1y-x1x1x1x1x1x1
- RELATED TARGETS:
 - EA: Azure-EA (11111111)
 - EA - Prod (x1x1x1x1-1a1a-x1x1-1a1a-x1x1x1x1x1)
 - EA - Test (x1x1x1x1-1a1a-x1x1-1a1a-x2x2x2x2x2)
 - Standalone
 - EATest (x1x1x1x1-1b1b-z1z1-1y1y-x1x1x1x1x1)

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

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

Validated: Jan 28, 2021 7:24:56 PM

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 reservations. 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 reservations 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 \(on page 146\)](#).

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 reservations. 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.

Empty Azure EA subscriptions that are not incurring any charges will not stitch with the Azure Billing target or the Azure EA target, and a discrepancy will occur in the offer ID of the subscription. Once the subscription incurs a charge, the stitching occurs and the subscription should correctly associate with the Azure Billing target with the correct offer ID.

- 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.

Reservations and Azure EA

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

To discover and manage reservations 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 reservations. Workload Optimization Manager combines this information to track:

- Utilization of reservations
- VMs covered by reservations
- VM costs (accounting for reservations)
- Purchase recommendations

NOTE:

This release of Workload Optimization Manager does not support discovery and management of reservations for Classic VMs, Classic Cloud Services, and Suppressed Core VMs.

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, reservations 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 actions to purchase reservations might not match associated costs you find in the Microsoft Pricing Calculator.

Workload Optimization Manager actions can recommend 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 on-prem workloads you migrated to Azure, Workload Optimization Manager recognizes Azure Hybrid Benefit (AHUB) savings for reservations 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	EMC XtremIO	HPE 3Par	NetApp	Nutanix	Pure
Storage	Volume (Regular, Thin, Meta)	Volume	Virtual Volume	Volume	Container	Volume
Disk Array	Disk Group or Thin Pool	XTremIO Cluster	CPG	Aggregate	Storage Pool	Shelf Array
Storage Controller	VMAX Array	XTremIO Cluster	Controller	Controller / Filer	Controller VM	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. The actions that Workload Optimization Manager can recommend or automate depend on the actual technology. 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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p>NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p>

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.
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 allocated CPU ▪ 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.

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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="margin-left: 40px;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Logical Pool	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the logical pool's capacity. ▪ Storage Provisioned The utilization of the logical pool's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the logical pool. ▪ Latency The utilization of latency on the logical pool.
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.
Storage Controller	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the storage controller's capacity.

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 \(on page 168\)](#).

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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p>NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.

Entity Type	Commodity
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 allocated CPU ▪ 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.

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

EMC ScaleIO Name	Workload Optimization Manager Entity
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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array

Entity Type	Commodity
	<ul style="list-style-type: none"> Latency The utilization of latency, computed from the latency of each device in the disk array.
Storage Controller	<ul style="list-style-type: none"> Storage Amount The utilization of the storage controller's capacity.

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.

VPLEX Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> Service Account
Action Execution	<ul style="list-style-type: none"> Admin account

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 \(on page 168\)](#).

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
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 Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p>NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> Storage Amount The utilization of the Disk Array's capacity. Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. 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> <ul style="list-style-type: none"> ▪ Latency <p>The utilization of latency, computed from the latency of each device in the disk array.</p>
Logical Pool	<ul style="list-style-type: none"> ▪ Storage Amount <p>The utilization of the logical pool's capacity.</p> <ul style="list-style-type: none"> ▪ Storage Provisioned <p>The utilization of the logical pool's capacity, including overprovisioning.</p> <ul style="list-style-type: none"> ▪ Storage Access Operations Per Second (IOPS) <p>The summation of the read and write access operations per second on the logical pool.</p> <ul style="list-style-type: none"> ▪ Latency <p>The utilization of latency on the logical pool.</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 allocated CPU</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>

IBM FlashSystem

IBM® FlashSystem family comprises all-Flash storage-array platforms. The platform delivers an active-active dual controller system. Assuming a pair of controller nodes, the active-active controller pair can process I/O for a specific volume through either node.

A FlashSystem control enclosure contains multiple NVMe-attached IBM FlashCore® Modules or other self-encryption NVMe-attached SSD drives. Each control enclosure contains two identical node canisters. A node canister provides host interfaces, management interfaces, and interfaces to the control enclosure. The pair of nodes within a single enclosure is known as an input/output (I/O) group.

NOTE:

For deployments that include mirrored volumes, Workload Optimization Manager discovers the Primary copy, and uses data from that for analysis and to generate actions. It does not discover the Secondary copy.

Prerequisites

- A service user account on the FlashSystem controller

You must configure a user account on the FlashSystem with a `Monitor` role. This user account must be dedicated to a single Workload Optimization Manager instance. To manage one FlashSystem with multiple Workload Optimization Manager instances, you must create a different user account for each Workload Optimization Manager instance.

- Spectrum Virtualize REST version 8.3.1.2 or later (8.4.2.0 or later recommended)

To take advantage of significant improvements to the Spectrum Virtualize interface, we recommend that you use Spectrum Virtualize version 8.4.2.0 or later.

IBM FlashSystem Permissions

Workload Optimization Manager Functionality	Required Permissions
Monitoring	<ul style="list-style-type: none"> ▪ Service Account with a <code>Monitor</code> role

Considerations

When you configure a FlashSystems target, you should consider:

- FlashSystem devices under IBM SVC management are not currently supported.
- Workload Optimization Manager does not discover mirrored volumes or child pools.
- Clusters must consist of a single enclosure.

Adding IBM FlashSystem Targets

For FlashSystem targets, select the **Storage > IBM FlashSystem** option on the Target Configuration page. Then provide the following:

- Address
The host name or Management IP address of the FlashSystem Cluster.
- Username
The username for a dedicated Workload Optimization Manager user account on the FlashSystem instance.
- Password
The password for the given user account.

Supply Chain

IBM FlashSystem Cluster targets add Storage Controller, Disk Array, and Logical Pool entities to the supply chain. Logical Pool entities provide Storage entities, which the supply chain stitches to Hypervisor storage (for example, vCenter datastores).

The FlashSystem components that Workload Optimization Manager discovers map to supply chain entities as follows:

FlashSystem Name	Supply Chain Entity	Description
Volume (VDisk)	Storage	A logical disk that the system presents to attached hosts. You create Volumes from the extents that are available in the pool.
Pool (MDiskGroup)	Logical Pool	An allocated amount of capacity that jointly contains all of the data for a specified set of volumes. All MDisks in

FlashSystem Name	Supply Chain Entity	Description
		a pool are split into extents of the same size.
MDisk (Array)	Disk Array	A logical unit of physical storage. MDisks are not visible to host systems.
System (Enclosure or Controller)	Storage Controller	A FlashSystem control enclosure.

Supported Actions

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

Entity	Actions
Storage Controller	<ul style="list-style-type: none"> ▪ Disabled by default <ul style="list-style-type: none"> – Provision
Disk Array	<ul style="list-style-type: none"> ▪ Disabled by default <ul style="list-style-type: none"> – Resize – Move – Provision – Start – Suspend
Logical Pool	<ul style="list-style-type: none"> ▪ Recommend by default <ul style="list-style-type: none"> – Resize (add or remove MDisk) ▪ Disabled by default <ul style="list-style-type: none"> – Move – Provision – Start – Suspend
Storage	<ul style="list-style-type: none"> ▪ Manual by default <ul style="list-style-type: none"> – Delete ▪ Recommend by default <ul style="list-style-type: none"> – Resize – Move – Provision – Start ▪ Disabled by default <ul style="list-style-type: none"> – Suspend
Virtual Machine	<ul style="list-style-type: none"> ▪ Recommend by default <ul style="list-style-type: none"> – Storage vMotion

Monitored Resources

To discover and monitor resources, Workload Optimization Manager uses the Spectrum Virtualize REST interface, and it securely retrieves `iostats` files from the configuration node.

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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="text-align: center;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Logical Pool	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the logical pool's capacity. ▪ Storage Provisioned The utilization of the logical pool's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the logical pool. ▪ Latency The utilization of latency on the logical pool.
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.
Storage Controller	<ul style="list-style-type: none"> ▪ CPU The utilization of the Storage Controller's allocated CPU

Entity Type	Commodity
	<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>

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 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 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 \(on page 187\)](#)

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 cluster management server.
- 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 \(on page 168\)](#).

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	Provision, Resize Up
Disk Array		Resize Up, Move, Provision
Storage Controller		Provision

Note that 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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore

Entity Type	Commodity
	<p>NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.
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 allocated CPU ▪ 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.

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 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 \(on page 168\)](#).

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 ▪ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore <p style="text-align: center;">NOTE: When it generates actions, Workload Optimization Manager does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities.</p> <ul style="list-style-type: none"> ▪ Latency The utilization of latency on the datastore
Disk Array	<ul style="list-style-type: none"> ▪ Storage Amount The utilization of the Disk Array's capacity. ▪ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. ▪ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array ▪ Latency The utilization of latency, computed from the latency of each device in the disk array.

Entity Type	Commodity
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 allocated CPU▪ 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.



Virtual Desktop Infrastructure Targets

A virtual desktop infrastructure target is a service that delivers digital workspaces to multiple application users through a single portal. When you connect Workload Optimization Manager to VDI targets, it monitors the performance and resource consumption of the business users utilizing the virtual desktops and the virtual desktop infrastructure itself. Workload Optimization Manager then performs Application Resource Management (ARM) using the enhanced data to assure application performance while also utilizing resources as efficiently as possible.

As part of this process, Workload Optimization Manager will use this information to provide deeper insight into the state of the hardware, and information related to the users and virtual desktops running in your environment.

Monitored Resources

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

Entity Type	Commodity
Business User	<ul style="list-style-type: none">▪ ImageCPU CPU utilization, as a percentage of CPU capacity for the user's desktop image or images.▪ ImageMem Memory utilization, as a percentage of Memory capacity for the user's desktop image or images.▪ ImageStorage Storage utilization, as a percentage of storage capacity for the user's desktop image or images.
Desktop Pool	<ul style="list-style-type: none">▪ Pool CPU The CPU available to the pool that is in use by active sessions.▪ Pool Memory The memory available to the pool that is in use by active sessions.▪ Pool Storage The storage capacity available to the pool that is in use by active sessions.▪ Active Sessions How many active sessions are on the pool as a percentage of the pool's capacity as defined in the Workload Optimization Manager policy.

Entity Type	Commodity
	<ul style="list-style-type: none"> ▪ Total Sessions How many active and disconnected (non-terminated) sessions are on the pool, as a percentage of the pool's capacity.

Horizon

VMware Horizon provides end users access to all of their virtual desktops, applications, and online services through a single digital workspace.

With the information provided by Horizon, Workload Optimization Manager efficiently utilizes the VDI infrastructure to avoid inappropriately placed or sized desktops, ensure users are matched to desktop images that meet their demand, while gaining valuable insight into the interconnected nature of your environment.

Prerequisites

- Horizon 7.0 or higher
- The version of the vCenter target associated to the Horizon target must be 5.0+

Adding Horizon Targets

NOTE:

For each Horizon target, you must add only one connection server servicing that site. Workload Optimization Manager will use the default port of 443. Adding multiple connection servers for the same site will result in inconsistent and/or duplicate actions.

To add a Horizon target, select the **Virtual Desktop Infrastructure > Horizon** option on the Target Configuration page and provide the following information:

- **Address**
The IP address of the Horizon connection server.
- **Username**
Username for the user account Workload Optimization Manager will use to connect to the Horizon target. This account must be an Administrator (Read-Only) account.
- **Password**
Password for the user account Workload Optimization Manager will use to connect to the Horizon target.
- **vCenter Username**
Username for the user account Workload Optimization Manager will use to connect to the vCenter target related to Horizon. Leave this blank if the Horizon and vCenter credentials are identical.
- **vCenter Password**
Password for the user account Workload Optimization Manager will use to connect to the vCenter target related to Horizon. Leave this blank if the Horizon and vCenter credentials are identical.
- **LDAP Server**
The address of the specific LDAP server Workload Optimization Manager will use in conjunction with the Horizon target. By default, Workload Optimization Manager will do a lookup based on the domain entered in the Username and select a healthy DNS server from the response.
- **LDAP Server Port**
The port Workload Optimization Manager will use to connect to the specified LDAP server. By default, Workload Optimization Manager will use port 389 for unsecure connections, and port 636 for secure connections.

- **Secure LDAP Connection**
When checked, Workload Optimization Manager will use SSL to communicate to the LDAP Server
- **LDAP Username**
The username Workload Optimization Manager will use to connect to the LDAP Server, in the user principal name format to authenticate in the directory service through Kerberos. For example, `user.name@domain.com`
- **LDAP Password**
The password Workload Optimization Manager will use to connect to the LDAP Server

Actions

Workload Optimization Manager recommends actions for the various entities of the Horizon target as follows:

Entity Type	Action
Desktop Pool	Workload Optimization Manager does not recommend actions to perform on a desktop pool. Instead, it recommends actions to perform on the business users running active sessions in the pool.
View Pod	Workload Optimization Manager does not recommend actions to perform on a view pod. Instead, it recommends actions to perform on the business users running active sessions in the view pod.
Business User	<ul style="list-style-type: none"> ▪ Move Business User due to image congestion ▪ Move Business User due to desktop pool congestion

Under some circumstances, you can have Business Users who need larger desktop images. This appears as users with high utilization of the image resources. In this case, Workload Optimization Manager can recommend moving the Business Users to a different desktop pool that serves up larger images.

To support moving Business Users, you must create a placement policy that merges desktop pools. Be sure to merge only desktop pools that are *similarly configured* – they should run the same operating system and applications, and differ only in allocated memory and/or CPU.

To merge desktop pools:

1. Create a new placement policy.
2. Choose **Merge** as the policy type.
3. For the consumer type to merge, choose **Desktop Pool**.
4. Choose the pools that you want to merge.
5. Save the policy.

For more information, see "Creating Placement Policies" in the *User Guide*.

Monitored Resources

Workload Optimization Manager monitors the following commodities of the Horizon target:

Entity Type	Commodity
Business User	<ul style="list-style-type: none"> ▪ ImageCPU CPU utilization, as a percentage of CPU capacity for the user's desktop image or images. ▪ ImageMem

Entity Type	Commodity
	<p>Memory utilization, as a percentage of Memory capacity for the user's desktop image or images.</p> <ul style="list-style-type: none"> ▪ ImageStorage <p>Storage utilization, as a percentage of storage capacity for the user's desktop image or images.</p>
Desktop Pool	<ul style="list-style-type: none"> ▪ Pool CPU The CPU available to the pool that is in use by active sessions. ▪ Pool Memory The memory available to the pool that is in use by active sessions. ▪ Pool Storage The storage capacity available to the pool that is in use by active sessions. ▪ Active Sessions How many active sessions are on the pool as a percentage of the pool's capacity as defined in the Workload Optimization Manager policy. ▪ Total Sessions How many active and disconnected (non-terminated) sessions are on the pool, as a percentage of the pool's capacity.
View Pod	<ul style="list-style-type: none"> ▪ Active Sessions How many active sessions are on the pool as a percentage of the pool's capacity as defined in the Workload Optimization Manager policy. ▪ Total Sessions How many active and disconnected (non-terminated) sessions are on the pool, as a percentage of the pool's capacity.



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

We highly recommend enabling collection of memory metrics in your environment. Enabling memory metrics allows Workload Optimization Manager to generate Scale actions to optimize VM memory usage. For Workload Optimization Manager to collect memory metrics, you must enable the collection of these metrics on the VMs in your environment.

Some of the steps to do this are different depending on whether your VM is running a Linux or Windows OS. To enable memory statistics on your VMs, you must meet the following requirements:

- The VM image must have an SSM agent installed
 - Linux VMs:
By default, Linux AMIs dated 2017.09 and later include an installed SSM Agent.
 - Windows VMs:
You must install the SSM agent on the VMs. For more information, see [Working with 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 from Workload Optimization Manager

For Workload Optimization Manager to access the memory data, the account that it uses to connect to the AWS target must include the correct permissions. If you configured the AWS target via an AWS key (not an IAM role), then you must include the permissions as specified in the section for configuring an AWS target.

If you use an IAM role for the Workload Optimization Manager connection, then that role must include the following as a minimum:

- `AmazonEC2ReadOnlyAccess`
- `AmazonS3ReadOnlyAccess`
- `AmazonRDSReadOnlyAccess`

To set up the collection of memory statistics for your VMs:

1. Attach an IAM role to each VM instance.

Each EC2 instance must have an attached IAM role that grants CloudWatch access. To grant that access, we suggest you include the `AmazonSSMFullAccess` policy in the role.

Use the AWS System Manager to attach the necessary roles to your VMs.

NOTE:

If you want to grant the role lesser access, you can use the `AmazonEC2RoleforSSM` policy. This is a custom policy that allows the action `ssm:GetParameter` to access the resource, `arn:aws:ssm:*:*:parameter/*`.

2. Install the CloudWatch agent on your Linux VMs.

Navigate to the AWS System Manager service for the account and region that you want to configure. In the service, navigate to the **Run Command** screen and set up the **AWS-ConfigureAWSPackage** command to install **AmazonCloudWatchAgent** on your VMs. For more information, please see the AWS documentation.

3. Create configuration data for the CloudWatch agent.

The configuration data is a JSON object that you will add to as a parameter to the Parameter Store. The object must include the following, depending on whether it's for a Linux or a Windows VM instance.

- **Linux Configuration:**

```
{
  "agent": {
    "metrics_collection_interval": 60,
    "logfile": "/opt/aws/amazon-cloudwatch-agent/logs/amazon-cloudwatch-agent.log"
  },
  "metrics": {
    "namespace": "custom",
    "metrics_collected": {
      "mem": {
        "measurement": [
          {
            "name": "mem_available", "rename": "MemoryAvailable", "unit": "Bytes"
          }
        ]
      }
    },
    "append_dimensions": {
      "AutoScalingGroupName": "${aws:AutoScalingGroupName}",
      "ImageId": "${aws:ImageId}",
      "InstanceId": "${aws:InstanceId}",
      "InstanceType": "${aws:InstanceType}"
    }
  }
}
```

- **Windows Configuration:**

```
{
  "metrics": {
    "namespace": "Windows System",
    "append_dimensions": {
      "InstanceId": "${aws:InstanceId}"
    },
    "aggregation_dimensions": [ ["InstanceId"] ],
    "metrics_collected": {
      "Memory": {
        "measurement": [
          { "name": "Available Bytes", "rename": "MemoryAvailable", "unit": "Bytes" }
        ],
        "metrics_collection_interval": 60
      },
      "Paging File": {
        "measurement": [
          { "name": "% Usage", "rename": "paging_used" }
        ],
        "metrics_collection_interval": 60,
        "resources": [
          "*"
        ]
      }
    }
  }
}
```

Note that you can configure optional parameters for the CW Namespace and region. However, if you configure more metrics for CloudWatch to collect, these metrics do not affect Workload Optimization Manager analysis and they do not show up in the user interface. Workload Optimization Manager only tracks Used Memory statistics.

4. Create a parameter store.
 - a. Create a parameter.

In the AWS System Manager, navigate to **Parameter Store** and create a parameter. Copy and paste the JSON agent configuration (created above) into the parameter **Value** field.
 - b. Name the parameter

For example, `AmazonCloudWatch-MyMemoryParam`. You can use a different name, but per the Amazon documentation, the name *must* begin with `AmazonCloudWatch`. For more information, see [Store the CloudwatchConfig File in Parameter Store](#).

You must remember this parameter name.
 - c. Click to create the parameter.

5. Deploy the CloudWatch parameter to your VMs.

In the AWS System Manager, navigate to the **Run Command** screen to configure and run the **AmazonCloudWatch-ManageAgent** command. The configuration should include:

- **Action:** `configure`
- **Mode:** `ec2`
- **Optional Configuration Source:** `ssm`
- **Optional Configuration Location:** Give the name of the parameter that you created above.
- **Optional Restart:** `yes` (this restarts the CloudWatch Agent, not the VM instance)
- **Targets:** The VMs that you will deploy the CloudWatch configuration to

When the command is configured, run it. This configures collection of memory metrics for your instances.

6. Verify that you are collecting Memory metrics for your instances.

Navigate to the CloudWatch page, and display **Metrics** in the **CWAgent** namespace. Then inspect the instances by ID to verify that you can see `MemoryAvailable` metrics.

Memory Metrics Collection: Azure

We highly recommend enabling collection of memory metrics in your environment. Enabling memory metrics allows Workload Optimization Manager to generate Scale actions to optimize VM memory usage. For Workload Optimization Manager to collect memory metrics, you must enable the collection of these metrics on the VMs in your environment.

Workload Optimization Manager fetches memory metrics for Azure VMs using the following mechanisms, in order:

1. **Log Analytics Workspace - InsightsMetrics table** - (preferred method, supported by Azure) This mechanism is used when the VM is configured with the Azure Monitor Agent (AMA). This agent collects Guest OS Memory metrics and sends it to the Log Analytics Workspace `InsightsMetrics` table. Workload Optimization Manager then fetches these Available Memory metrics by querying this table.

For more information on the Azure Monitor Agent, see the [Azure documentation](#).

2. **Log Analytics Workspace - Perf table** - This mechanism is used when the VM is configured with the Log Analytics Agents (OMS/MMA) and the VM is connected to a Log Analytics Workspace. The Log Analytics Agents send the Guest OS Memory metrics to the Log Analytics Workspace `Perf` table. Workload Optimization Manager then fetches these Available Memory metrics by querying this table.

The Log Analytics Agents (OMS/MMA) will reach end of support by August 2024. For information on the recommended migration to the new Azure Monitor Agent (AMA), see the [Azure documentation](#).

3. **Storage Account SDK Based Diagnostics Monitoring** - This legacy VM memory collection mechanism installs an IaaS Diagnostics Agent on the VM, which sends guest OS level memory metrics to a designated Storage Account table. During discovery process, Workload Optimization Manager collects VM memory from these tables. If this mechanism is being used, Workload Optimization Manager needs relevant `StorageAccount` read permissions so that the memory metrics can be read from these tables.

4. **Metrics REST API** - If memory could not be collected for a VM using the previously-mentioned mechanisms, Workload Optimization Manager tries to fetch host OS based memory for the VM via REST API calls, which is the same mechanism used for collecting other VM metrics such as Percentage CPU usage. Workload Optimization Manager collects VM memory and tries to generate memory-based scaling actions where appropriate.

NOTE:

Azure provides Available Memory Bytes via REST API in Preview mode, or through the Azure Portal UI under the VM's **Monitoring > Metrics** section.

Azure Memory Source Group Discovery

Azure Memory Source groups reflect the source from where VM memory is being fetched.

In the Workload Optimization Manager user interface, to search for a specific Azure Memory Source group, choose **Groups** in the Search page and enter `memory` in the search field. Each group contains the list of VMs for which memory was fetched via the specified mechanism. These are dynamic groups that get continuously updated based on each discovery.

The Azure Memory Monitoring Unavailable group contains VMs for which memory could not be fetched using any of the collection mechanisms. This group may contain VMs that are not in a `RUNNING` power state, and metrics data is often not available for these VMs. Memory-related scaling actions are not generated for VMs in this group; however, other non-memory Scale actions may be generated.

The screenshot shows the Azure portal search interface. On the left is a navigation menu with 'Groups' highlighted. The main area has a search bar containing 'memory' and an 'ADD FILTER' button. Below the search bar is a table of search results:

			NAME	↑
Azure Memory Monitoring Available (Log Analytics InsightsMetrics)	78	Static	>	
<small>Cloud</small>	<small>Virtual Machines</small>			
Azure Memory Monitoring Available (Log Analytics Perf)	2	Static	>	
<small>Cloud</small>	<small>Virtual Machines</small>			
Azure Memory Monitoring Available (Metrics API)	55	Static	>	
<small>Cloud</small>	<small>Virtual Machines</small>			
Azure Memory Monitoring Available (Storage SDK Monitoring)	6	Static	>	
<small>Cloud</small>	<small>Virtual Machines</small>			
Azure Memory Monitoring Unavailable	52	Static	>	
<small>Cloud</small>	<small>Virtual Machines</small>			

GCP Target Service Account

This appendix guides you through the steps to create a valid service account that Workload Optimization Manager can use to connect with a GCP operational target. The access you grant to this service account determine the access Workload Optimization Manager has to discover and manage entities in your GCP infrastructure. To create this account you will:

- Create a service account in GCP
- Generate key file for the service account
- Create custom roles for the project(s) you want Workload Optimization Manager to manage
- Add the custom roles to the GCP Service Account, plus the predefined Billing Account Viewer role

Create a Service Account in GCP

Open a `gcloud shell` session in the project that you want to host the new Service Account. In `gcloud shell`, execute the following command, where:

- `$SERVICE_ACCOUNT_ID` is the unique alphanumeric ID that you assign to this service account
- `$SERVICE_ACCOUNT_DISPLAY_NAME` is the display name that you want for the service account

```
$ gcloud iam service-accounts create $SERVICE_ACCOUNT_ID \
  --display-name="$SERVICE_ACCOUNT_DISPLAY_NAME" \
  --format=text --quiet
```

When you execute the command, the shell should display the following output:

```
displayName: <SERVICE_ACCOUNT_DISPLAY_NAME>
email: <SERVICE_ACCOUNT_EMAIL>
etag: MDEwMjE5MjA=
name: <SERVICE_ACCOUNT_NAME>
projectId: <SA_DEFAULT_PROJECT>
uniqueId: 102200949905427524050
```

Note the following fields in the output for later use:

- `email`
- `name`
- `projectId`

NOTE:

The `projectId` field identifies the project that hosts this Service Account. You should note that project for future reference, in case you want to review or edit the Service Account.

Generate the Key File

To generate the key file, execute the following command in the `gcloud shell`:

```
$ gcloud iam service-accounts keys create $SA_KEY_FILE_NAME \
  --iam-account=$SERVICE_ACCOUNT_EMAIL
```

Be sure to save the key file – you need it when you configure the GCP Target in Workload Optimization Manager.

Create Custom Roles

You must create two different roles for access to your GCP projects and organization, and you can create a third role for permission to execute actions in projects.

NOTE:

To target specific folders, define a custom role at the organization level. It is not possible to define custom roles at the folder level.

1. Project Access Role:

In a location you can access from your `gcloud shell` session, create a file named `CWOMSaAccessProject.yaml`. Edit the file to have the following content:

```

title: "CWOM Role: Min Access - Project"
description: "Minimal Required Permissions for CWOM to manage the GCP Project"
stage: "ALPHA"
includedPermissions:
# discovery
- resourcemanager.projects.get
- compute.regions.list
- compute.zones.list
- compute.machineTypes.list
- compute.machineTypes.get
- compute.disks.list
- compute.disks.get
- compute.diskTypes.list
- compute.instances.list
- compute.instances.get
- compute.instanceGroupManagers.list
- compute.instanceGroupManagers.get
# CUD
- compute.commitments.list
# Metrics Monitoring
- logging.views.list
- logging.views.get
- monitoring.services.get
- monitoring.services.list
- monitoring.timeSeries.list
- serviceusage.services.get

```

Then execute the following command to create the custom role in your organization, where you substitute `CWOMSaAccessProject.yaml` with the path to your yaml file:

```

$ gcloud iam roles create $CUSTOM_ROLE_NAME \
  --project=$PROJECT_ID \
  --file=CWOMSaAccessProject.yaml

```

In the resulting console output, note the role name. You need that name to add the role to your service account.

2. Organization Access Role:

In a location you can access from your `gcloud` shell session, create a file named `CWOMSaAccessOrg.yaml`. Edit the file to have the following content:

```

title: "CWOM Role: Access - Organization"
description: "Minimal Required Permissions for CWOM to access the GCP Organization"
stage: "ALPHA"
includedPermissions:
# Organization Structure View
- resourcemanager.organizations.get
- resourcemanager.projects.list
- resourcemanager.projects.get
- resourcemanager.folders.list
- resourcemanager.folders.get
- billing.resourceAssociations.list

```

Then execute the following command to create the custom role in your organization, where you substitute `CWOMSaAccessOrg.yaml` with the path to your yaml file:

```

$ gcloud iam roles create $CUSTOM_ROLE_NAME \
  --organization=$ORGANIZATION_ID \
  --file=CWOMSaAccessOrg.yaml

```

In the resulting console output, note the role name. You need that name to add the role to your service account.

3. Project Action Role:

Create this role if you want Workload Optimization Manager to execute actions in your GCP environment.

In a location you can access from your `gcloud` shell session, create a file named `CWOMSaProjectAction.yaml`. Edit the file to have the following content:

```
title: "CWOM Role: Project Action Execution"
  description: "Grant CWOM to execute actions in the GCP Project"
stage: "ALPHA"
includedPermissions:
# Action Execution
- compute.globalOperations.get
- compute.instances.setMachineType
- compute.instances.start
- compute.instances.stop
- compute.disks.delete
- compute.regionOperations.get
- compute.zoneOperations.get
```

Then execute the following command to create the custom role in your organization, where you substitute `CWOMSaProjectAction.yaml` with the path to your yaml file:

```
$ gcloud iam roles create $CUSTOM_ROLE_NAME \
  --project=$PROJECT_ID \
  --file=CWOMSaProjectAction.yaml
```

In the resulting console output, note the role name. You need that name to add the role to your service account.

Add Custom Roles to the Service Account

You will now add the three roles you created to your Service Account, plus the predefined Billing Viewer role.

1. (Required) Add the predefined Billing Viewer role.

In the GCP Console, go to **Billing > Account Management**. This page should display the billing account and project you are going to use as the GCP Billing target in Workload Optimization Manager.

Add the **Billing Account Viewer** role to the Service Account

2. (Required) Add the *Organization Access* role that you created for the Service Account.

In the `gcloud` console session, execute this command, where `$ROLE_NAME` is the role name you noted when you created the Organization Access role:

```
$ gcloud projects add-iam-policy-binding $PROJECT_ID \
  --member=serviceAccount:$SERVICE_ACCOUNT_EMAIL --role=$ROLE_NAME
```

3. (Required) Add the *Project Access* role that you created for the Service Account.

In the `gcloud` console session, execute this command, where `$ROLE_NAME` is the role name you noted when you created the Project Access role:

```
$ gcloud projects add-iam-policy-binding $PROJECT_ID \
  --member=serviceAccount:$SERVICE_ACCOUNT_EMAIL --role=$ROLE_NAME
```

4. (Optional) Add the *Project Action* role that you created for the Service Account.

You only need to add this role if you want Workload Optimization Manager to execute actions in your GCP project.

In the `gcloud` console session, execute this command, where `$ROLE_NAME` is the role name you noted when you created the Project Action role:

```
$ gcloud projects add-iam-policy-binding $PROJECT_ID \
  --member=serviceAccount:$SERVICE_ACCOUNT_EMAIL --role=$ROLE_NAME
```

You now have created a Service Account that you can use to configure your GCP operational target.

GCP Billing Target Service Account

This appendix guides you through the steps to create a valid service account that Workload Optimization Manager can use to connect with a GCP Billing target. To create this account you will:

- Create a service account in GCP
- Generate key file for the service account
- Create a custom role in the project(s) you want Workload Optimization Manager to manage
- Add the created custom role to the GCP Service Account, plus the predefined Billing Account Viewer role

Create a Service Account in GCP

Open a `gcloud shell` session in the project that you want to host the new Service Account. In `gcloud shell`, execute the following command, where:

- `$SERVICE_ACCOUNT_ID` is the unique alphanumeric ID that you assign to this service account
- `$SERVICE_ACCOUNT_DISPLAY_NAME` is the display name that you want for the service account

```
$ gcloud iam service-accounts create $SERVICE_ACCOUNT_ID \
  --display-name="$SERVICE_ACCOUNT_DISPLAY_NAME" \
  --format=text --quiet
```

When you execute the command, the shell should display the following output:

```
displayName: <SERVICE_ACCOUNT_DISPLAY_NAME>
email: <SERVICE_ACCOUNT_EMAIL>
etag: MDEwMjE5MjA=
name: <SERVICE_ACCOUNT_NAME>
projectId: <SA_DEFAULT_PROJECT>
uniqueId: 102200949905427524050
```

Note the following fields in the output for later use:

- `email`
- `name`
- `projectId`

NOTE:

The `projectId` field identifies the project that hosts this Service Account. You should note that project for future reference, in case you want to review or edit the Service Account.

Generate the Key File

To generate the key file, execute the following command in the `gcloud shell`:

```
$ gcloud iam service-accounts keys create $SA_KEY_FILE_NAME \
  --iam-account=$SERVICE_ACCOUNT_EMAIL
```

Be sure to save the key file – you need it when you configure the GCP Billing target in Workload Optimization Manager

Create a Custom Billing Role

In a location you can access from your `gcloud shell` session, create a file named `CWOMSaBilling.yaml`. Edit the file to have the following content:

```

title: "CWOM Billing Data Viewer Role"
description: "Minimal Required Permissions for CWOM to view billed cost and pricing stored in the GCP Project"
stage: "ALPHA"
includedPermissions:
- bigquery.tables.get
- bigquery.tables.getData
- bigquery.tables.list
- bigquery.jobs.create
# derived cost probe will need
- compute.regions.list
- compute.zones.list
- compute.commitments.list
- compute.diskTypes.list
- compute.machineTypes.list

```

Then execute the following command to create the custom role in your organization, where you substitute `CWOMSaBilling.yaml` with the path to your yaml file:

```

$ gcloud iam roles create $CUSTOM_ROLE_NAME \
  --project=$PROJECT_ID \
  --file=CWOMSaBilling.yaml

```

In the resulting console output, note the role name. You need that name to add the role to your service account.

Add the Custom Roles to the Service Account

You will now add the Custom Billing role you created to your Service Account, plus the predefined Billing Viewer role.

1. (Required) Add the predefined Billing Viewer role.

In the GCP Console, go to **Billing > Account Management**. This page should display the billing account and project you are going to use as the GCP Billing target in Workload Optimization Manager.

Add the **Billing Account Viewer** role to the Service Account

2. (Required) Add the *Custom Billing* role that you created for the Service Account.

In the `gcloud` console session, execute this command, where `$ROLE_NAME` is the role name you noted when you created the Custom Billing role:

```

$ gcloud projects add-iam-policy-binding $PROJECT_ID \
  --member=serviceAccount:$SERVICE_ACCOUNT_EMAIL --role=$ROLE_NAME

```

You now have created a Service Account that you can use to configure your GCP Billing target.

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 \(on page 207\)](#)
- [EnablingWinRMViaPowerShell \(on page 208\)](#)

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 re-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 your intended targets, Workload Optimization Manager must be able to reach the targets via the following ports:

Applications and Databases

Target	Ports
AppDynamics	80, 443
JBoss	9990
MySQL	3306
Oracle	1521
SQL Server	1434, 1433
Tomcat	1009
WebLogic	7001
WebSphere	8880

Fabric and Network

Target	Ports
Cisco UCS	80, 443

Hyperconverged

Target	Ports
Nutanix	9440

Hypervisors

Target	Ports
Hyper-V	5985, 5986
Hyper-V (Kubernetes)	88
vCenter (Monitoring)	443
vCenter (Tags)	10443
vCenter (Kubernetes)	88

Private Cloud

Target	Ports
Microsoft VMM	5985, 5986

Storage

Target	Ports
EMC VMAX	5988, 5989
HPE 3PAR	5988, 5989, 8080
IBM FlashSystem	7443
NetApp	80, 443
Pure Storage	80, 443

Virtual Desktop Infrastructure

Target	Ports
VMware Horizon	443

AWS Target IAM Role Requirements

This appendix guides you through the steps to configure the Workload Optimization Manager AWS Mediation pods to leverage IAM roles. To do so, you must leverage the ability to provide fine grained IAM role support through a Service Account, and Kubernetes cluster configurations that support an OIDC provider and webhook method.

Prerequisites

Follow the steps below to configure the AWS Mediation pods to leverage IAM roles:

- Follow the instructions provided by Amazon EKS, Red Hat OpenShift, or GKE to make sure you have the required configurations to support the Web Identity provider method, leveraging the AWS webhook and an OIDC provider.
 - [Introducing EKS Granular IAM Roles for Pods via Service Accounts and OIDC Providers](#)
 - [Overview of setting up IAM Role for EKS WebIdentity Provider](#)
 - [Creating a Service Account with an IAM Role: EKS example](#)
 - [Red Hat OpenShift support for Granular IAM Roles](#)
 - GKE and IAM roles described in this [blog](#) using [gtoken](#) to inject the token into the pod
 - [AWS IAM Role Permissions and Trust Relationships](#)

NOTE:

Follow the instructions from AWS in the links above to set up your cluster's OIDC provider as an IAM Identity Provider (Web Identity Provider - OIDC URL) in the AWS account that you will be targeting.

2. Configure a Kubernetes Service Account in the Workload Optimization Manager namespace that will assume an IAM role.
 - (Best Practice) Manually create a separate Service Account for the AWS Mediation pods to use. You must then modify the Custom Resource YAML to specify this Service Account to the AWS Mediation components. For example:

```
spec:
  mediation-aws:
    serviceAccountName: t8c-iam-role
  mediation-awsbilling:
    serviceAccountName: t8c-iam-role
  mediation-awscost:
    serviceAccountName: t8c-iam-role
```

After you update the Custom Resource YAML, apply the updated Custom Resource and ensure the three Mediation pods restarted.

- Use the Workload Optimization Manager default Service Account with which the AWS Mediation pods are running.

NOTE:

Unless specified in the Workload Optimization Manager Custom Resource, the AWS Mediation pods will run with the default service account called "default" in the namespace. If you modify the default "default" account, there is nothing more to do except restart the AWS Mediation pods.

3. Configure the IAM role in AWS.
 - Workload Optimization Manager AWS IAM role requirements, including cross account access if required, are described [here](#). IAM policy definition must use the `sts:AssumeRoleWithWebIdentity` role.
 - Using the Web identity provider set up, and the Service Account you will use, update the Trust Relationships in the IAM role. See the instructions [here](#).
 - The IAM role for Workload Optimization Manager to target your AWS account requires the following policies:

```
AmazonRDSReadOnlyAccess
AmazonEC2ReadOnlyAccess
AmazonS3ReadOnlyAccess
AWSOrganizationsReadOnlyAccess
```

4. Annotate the Service Account you will use with the IAM role. See [this article](#) for an example.

```
apiVersion: v1
kind: ServiceAccount
#use the name of the SA that will contain the annotation
name: default
metadata:
  annotations:
    eks.amazonaws.com/role-arn: arn:aws:iam::<AWS_ACCOUNT_ID>:role/<IAM_ROLE_NAME>
```

5. Complete the Workload Optimization Manager AWS Account Target setup by going to the Workload Optimization Manager UI to configure the AWS Account targets using an IAM role. See [Amazon Web Services \(on page 124\)](#).