Cisco Intersight Workload Optimizer Target Configuration Guide

Contents

Target Configuration

Published on February 27, 2023

A target is a resource or workload management service in your virtual environment that Intersight Workload Optimizer has connected to. For example, a public cloud account on Amazon Web Services (AWS) can be a target, as can an on-prem datacenter managed by VMware vCenter Server. Cisco Intersight Workload Optimizer uses targets to monitor workloads and resources, and to execute actions in your environment. This guide describes how to configure targets to be claimed by your Intersight Workload Optimizer account.

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

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

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

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

The following Intersight Workload Optimizer targets are available for claiming in Intersight:

Category	Target Name	Minimum License Tier Required for Intersight Workload Optimizer	Intersight Assist Required
Cloud	Amazon Web Services	IWO Essentials	No
	Amazon Web Services Billing	IWO Essentials	No
	Google Cloud Platform	IWO Essentials	No
	Google Cloud Platform Billing	IWO Essentials	No
	Microsoft Azure Service Principal	IWO Essentials	No
	Microsoft Azure Billing	IWO Essentials	No

Category	Target Name	Minimum License Tier Required for Intersight Workload Optimizer	Intersight Assist Required
	Microsoft Azure Enterprise Agreement	IWO Essentials	No
Cloud Native	Kubernetes (IKS, CCP, OpenShift, EKS, AKS, GKE) ***Deployed on-premises***	IWO Advantage	Yes
	Kubernetes (OpenShift, EKS, AKS, GKE) ***SaaS or deployed on public cloud***	IWO Advantage	No
Applications and Databases	Apache Tomcat 7.x, 8.x, and 8.5.x ***Deployed on-premises***	IWO Advantage	Yes
	IBM WebSphere Application Server 8.5+ ***Deployed on-premises***	IWO Advantage	Yes
	JVM 6.0+ ***Deployed on-premises***	IWO Advantage	Yes
	Microsoft SQL Server 2012, 2014, 2016, 2017, and 2019 ***Deployed on-premises***	IWO Advantage	Yes
	MySQL Server 5.6.x and 5.7.x ***Deployed on-premises***	IWO Advantage	Yes
	Oracle 11g R2, 12c, 18c, and 19c ***Deployed on-premises***	IWO Advantage	Yes
Compute / Fabric	Cisco UCS Server (Standalone)	IWO Essentials	No
	Cisco UCS Domain (UCSM Managed)	IWO Essentials	No
	Cisco UCS Domain (Intersight Managed)	IWO Essentials	No
	HPE OneView 3.00.04	IWO Essentials	No
Guest OS Process / APM (Application Performance Management)	New Relic ***SaaS or deployed on public cloud***	IWO Premier	No
	Cisco AppDynamics 4.1+ ***Deployed on-premises***	IWO Advantage	Yes
	Cisco AppDynamics ***SaaS or deployed on public cloud***	IWO Advantage	No

Category	Target Name	Minimum License Tier Required for Intersight Workload Optimizer	Intersight Assist Required
	Dynatrace 1.1+ ***Deployed on-premises***	IWO Premier	Yes
	Dynatrace ***SaaS or deployed on public cloud***	IWO Premier	No
Hyperconverged	Cisco Hyperflex 3.5	IWO Essentials	No
	Nutanix Acropolis	IWO Essentials	Yes
Hypervisor	Microsoft Hyper-V 2008 R2, 2012/2012 R2, 2016, 2019	IWO Essentials	Yes
	VMware vCenter 6.0, 6.5, 6.7, and 7.0+	IWO Essentials	Yes
Change Management	ServiceNow	IWO Advantage	No
Storage	HPE 3PAR	IWO Essentials	Yes
	Dell EMC SC Series	IWO Essentials	Yes
	EMC VMAX using SMI-S 8.1+	IWO Essentials	Yes
	EMC ScaleIO 2.x and 3.x	IWO Essentials	Yes
	EMC VPLEX	IWO Essentials	Yes
	NetApp ONTAP 8.0+	IWO Essentials	Yes
	Pure Storage FlashArray	IWO Essentials	Yes
	EMC XtremeIO XMS 4.0+	IWO Essentials	Yes

Transport Layer Security Requirements

Intersight Workload Optimizer 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 Intersight Workload Optimizer 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.

Cloud Targets

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 Intersight Workload Optimizer 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. Intersight Workload Optimizer can analyze the performance of applications running on the public cloud, and provision more instances as demand requires. For a hybrid environment, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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. Intersight Workload Optimizer supports the following cloud technologies:

Supply Chain

For public clouds, Intersight Workload Optimizer 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 Intersight Workload Optimizer discovers the resources available to you through that account.

In order to discover RI utilization, you must provide Intersight Workload Optimizer with access to the S3 bucket that contains the AWS Cost and Usage report. Without this access, Intersight Workload Optimizer'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, Intersight Workload Optimizer will discover only those RIs that are purchased by accounts which have been added as targets.

Supported Regions – AWS

Intersight Workload Optimizer 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)	

Claiming AWS Targets

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

To add an AWS target, 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.

Whitelisting AWS Regions

While Intersight Workload Optimizer 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 Intersight Workload Optimizer from reaching all AWS regions. For example, if you operate Intersight Workload Optimizer 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 Intersight Workload Optimizer to discover.

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

Cost and Usage Report

In order for Intersight Workload Optimizer 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 AWS environment. With memory metrics, Intersight Workload Optimizer can generate actions that not only boost performance but also maximize your savings.

For Intersight Workload Optimizer to collect memory statistics in AWS, you must set up CloudWatch to enable the collection of these statistics on the VMs in your environment. For more information, see Enabling Collection of Memory Statistics: AWS *(on page 146)*.

AWS Permissions

The account for an AWS target must have permissions that enable Intersight Workload Optimizer 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 Intersight Workload Optimizer 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)	■ AmazonEC2ReadOnlyAccess ■ AmazonS3ReadOnlyAccess ■ AmazonRDSReadOnlyAccess ■ AWSConfigRoleForOrganizations (only required for consolidated billing with the master account)
Execute Actions	■ AmazonEC2FullAccess ■ AmazonS3ReadOnlyAccess ■ AmazonRDSFullAccess ■ AWSConfigRoleForOrganizations (only required for consolidated billing with the master account)

Minimum Permissions

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

Intersight Workload Optimizer Functionality	Required Permissions
Monitoring	■ autoscaling:DescribeAutoScalingGroups ■ cloudwatch:GetMetricData ■ cloudwatch:GetMetricStatistics

Intersight Workload Optimizer Functionality	Required Permissions
	■ 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 ■ rds:DescribeOrderableDBInstanceOptions ■ servicecatalog:SearchProducts
Action Execution	■ 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

Intersight Workload Optimizer Functionality	Required Permissions
	■ servicecatalog:DescribeProduct ■ servicecatalog:ProvisionProduct ■ servicecatalog:DescribeRecord ■ servicecatalog:ListLaunchPaths
Savings Plans	■ savingsplans:DescribeSavingsPlans

To enable access to an AWS Billing target, the account must also include these permissions (see AWS Billing Targets *(on page 16)*):

Intersight Workload Optimizer Functionality	Required Permissions
Monitoring	■ ce:GetReservationUtilization ■ ce:GetSavingsPlansUtilizationDetails ■ ce:GetSavingsPlansUtilization ■ ce:GetSavingsPlansCoverage ■ 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 Intersight Workload Optimizer supports.

1. Navigate to **More > 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.

Intersight Workload Optimizer 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

Intersight Workload Optimizer 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.
  - **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 Intersight Workload Optimizer will recommend actions to purchase additional discounts.

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

Intersight Workload Optimizer monitors the following resources for the cloud target supply chain:

Entity Type	Commodity
Virtual Machine (AWS)	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) ■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS ■ Net Throughput Rate of message delivery over a port Measured in KB/s ■ Net Throughput Inbound Rate of message received over a port Measured in KB/s ■ Net Throughput Outbound Rate of message sent over a port Measured in KB/s ■ I/O Throughput The throughput to the underlying storage for the entity Measured in KB/s ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Database Server	■ Virtual Memory (VMem) The utilization of VMem allocated to the Database Server's instance type ■ Virtual CPU (VCPU) The utilization of VCPU allocated to the Database Server's instance type

Entity Type	Commodity
	■ Storage Amount The amount of Amazon EBS storage utilized by the Database Server ■ Storage Access IOPS utilized by the Database Server ■ DB Cache Hit Rate (if available) The percentage of database responses served through cache hits ■ Connections The number of connections to the Database Server
Volume	■ Storage Access The percentage of the volume's capacity for storage access operations (measured in IOPS) that is in use. ■ IO Throughput The percentage of the volume's capacity for IO throughput (measured in MB/s) that is in use. ■ IO Throughput Read The percentage of the volume's capacity for IO throughput Read (measured in MB/s) that is in use. ■ IO Throughput Write The percentage of the volume's capacity for IO throughput Write (measured in MB/s) that is in use.

AWS Billing Families

As you configure AWS targets, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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 Intersight Workload Optimizer, if you enable Billing Family Recognition, then you can see the billing family master and member accounts in the Targets user interface, and Intersight Workload Optimizer 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 Intersight Workload Optimizer credentials for an AWS role that includes the permission, `organizations:DescribeOrganization`. With that permission, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer cannot discover complete billing information for that family, and its analysis will be based on incomplete information.

Intersight Workload Optimizer displays member accounts that have been configured as targets in regular text. For members that Intersight Workload Optimizer discovers but have not been configured as targets, Intersight Workload Optimizer 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

  Intersight Workload Optimizer 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 Intersight Workload Optimizer includes the new member in its analysis of billing families. If the new member is not already a target, then Intersight Workload Optimizer lists the new member in grayed text.
- You can configure discounts per billing family

  Intersight Workload Optimizer 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 Intersight Workload Optimizer user interface as a master account, but it has no member accounts.

AWS Billing Targets

The Intersight Workload Optimizer 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 Intersight Workload Optimizer instance.

Cloud service providers can offer their own price lists, including special costs for services or discounts for workloads. However, Intersight Workload Optimizer does not discover these adjustments. For example, to reflect any discounted prices in the Intersight Workload Optimizer display and in Intersight Workload Optimizer analysis, you must manually configure those discounts. In Intersight Workload Optimizer, 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, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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

Claiming an AWS Billing Target

To add an AWS Billing target, 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.

AWS Billing Target Permissions

Intersight Workload Optimizer Functionality	Required Permissions
Monitoring	■ ce:GetReservationUtilization ■ ce:GetSavingsPlansUtilizationDetails ■ ce:GetSavingsPlansUtilization ■ ce:GetSavingsPlansCoverage ■ organizations:DescribeOrganization ■ organizations:ListAccounts ■ s3:GetBucketAcl ■ s3:GetObject ■ sts:AssumeRole ■ sts:AssumeRoleWithWebIdentity

Actions

Intersight Workload Optimizer 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

Intersight Workload Optimizer 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. Intersight Workload Optimizer gains access to this infrastructure through a GCP service account with the appropriate permissions.

When you add a service account as a target, Intersight Workload Optimizer 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 Intersight Workload Optimizer but can be validated like any other target.

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

Intersight Workload Optimizer 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, Intersight Workload Optimizer 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

Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 148)* and GCP Billing Target Service Account *(on page 152)*.

Intersight Workload Optimizer Functionality	Minimum Authorization Scopes	Minimum IAM Roles/Permissions
(Required) Resource discovery	*Project-level scopes and permissions*	
	One of the following scope sets: Set 1: ■ https://www.googleapis.com/auth/cloudplatformprojects.readonly ■ https://www.googleapis.com/auth/cloud-platform ■ https://www.googleapis.com/auth/cloud-billing.readonly Set 2: ■ https://www.googleapis.com/auth/cloudplatformprojects.readonly ■ https://www.googleapis.com/auth/compute ■ https://www.googleapis.com/auth/monitoring.read ■ https://www.googleapis.com/auth/cloud-billing.readonly	■ 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 ■ logging.views.list ■ logging.views.get ■ monitoring.services.get ■ monitoring.services.list ■ monitoring.timeSeries.list ■ serviceusage.services.get
	Organization-level scopes and permissions **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.	
	One of the following scope sets: Set 1: ■ https://www.googleapis.com/auth/cloud-platform ■ https://www.googleapis.com/auth/cloud-billing.readonly Set 2: ■ https://www.googleapis.com/auth/cloud-platform.read-only ■ https://www.googleapis.com/auth/iam.test	■ All project-level permissions ■ resourcemanager.organizations.get ■ resourcemanager.folders.get ■ resourcemanager.folders.list ■ resourcemanager.projects.list ■ resourcemanager.projects.get ■ billing.resourceAssociations.list

Intersight Workload Optimizer Functionality	Minimum Authorization Scopes	Minimum IAM Roles/Permissions
	■ https://www.googleapis.com/auth/cloud-billing.readonly Set 3: ■ https://www.googleapis.com/auth/cloudplatformorganizations.readonly ■ https://www.googleapis.com/auth/cloudplatformfolders.readonly ■ https://www.googleapis.com/auth/cloudplatformprojects.readonly ■ https://www.googleapis.com/auth/iam.test ■ https://www.googleapis.com/auth/cloud-billing.readonly	
(Required) Cost discovery Intersight Workload Optimizer discovers billing families, billed costs, negotiated pricing, and committed use discounts.	■ All project-level scopes ■ https://www.googleapis.com/auth/cloud-billing.readonly ■ One of the following scopes: – https://www.googleapis.com/auth/compute – https://www.googleapis.com/auth/cloud-platform	The service account should have the Billing Account Viewer role and the following permissions to the project that stores billing data: ■ All project-level permissions ■ billing.resourceAssociations.list ■ billing.accounts.list ■ compute.commitments.list 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: ■ bigquery.tables.get ■ bigquery.tables.getData ■ bigquery.tables.list ■ bigquery.jobs.create
(Optional) Action execution	■ All project-level scopes ■ One of the following scopes: – https://www.googleapis.com/auth/compute – https://www.googleapis.com/auth/cloud-platform	■ All project-level permissions ■ compute.instances.stop ■ compute.instances.setMachineType ■ compute.instances.start ■ compute.disks.delete ■ compute.zoneOperations.get ■ compute.regionOperations.get ■ compute.globalOperations.get

Enabling Collection of Memory Metrics

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

GCP collects these metrics via Ops Agent. In order for Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer, 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 Intersight Workload Optimizer. 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.

Adding Billing Accounts as Targets

To add a billing account target, specify the following:

- Target Name

  The name that identifies the target in Intersight Workload Optimizer. 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. Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer, 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**.

Cloud Instance Family Support

In the user interface, you can see the instance types that Intersight Workload Optimizer supports.

1. Navigate to **More > 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.

Intersight Workload Optimizer 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, Intersight Workload Optimizer updates the supply chain with the entities that it discovered. The following table describes the entity mapping between GCP and Intersight Workload Optimizer.

GCP	Intersight Workload Optimizer
Virtual Machine (VM) Instance	Virtual Machine (VM) **NOTE:** Intersight Workload Optimizer 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 Intersight Workload Optimizer (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 Intersight Workload Optimizer uses UTC, while GCP uses local time. However, costs shown in both places are correct and reliable.

Monitored Resources

Intersight Workload Optimizer monitors the following resources for GCP workloads:

Entity Type	Commodity
Virtual Machine	**NOTE:** Intersight Workload Optimizer does not monitor GCP machine types or CPU platforms that are currently in the preview/beta state. ■ Virtual Memory (vMem) capacity utilized by the VM In order for Intersight Workload Optimizer to retrieve vMem metrics, you must install and configure Ops Agent on each VM that it monitors. ■ Virtual CPU (vCPU) capacity utilized by the VM Intersight Workload Optimizer 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 Intersight Workload Optimizer calculates capacity or uses GCP-published capacity data, depending on the VM's machine type and disk. – Shared-core machine types share a physical core and are used for running small, non-resource intensive apps. • For shared-core machine types with *standard disks*, Intersight Workload Optimizer uses capacity data that GCP publishes here. Note that there

Entity Type	Commodity
	are no published values for e2-micro and e2-small, so Intersight Workload Optimizer assumes the e2-medium capacity for these machine types. • For shared-core machine types with *SSDs*, Intersight Workload Optimizer 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 *not* shared-core: • Intersight Workload Optimizer uses published capacity data and assumes that I/O block size is 16KB per I/O. • For machine types with persistent disks, Intersight Workload Optimizer assumes that the published capacity for the *SSD* disk type also applies to the *balanced* and *extreme* disk types. 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 *standard* disk type, capacity is the per-VM limit for the standard disk type. ■ Net throughput (inbound and outbound) for a VM
Volume	Currently, Intersight Workload Optimizer 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, Intersight Workload Optimizer analyzes resource utilization percentiles and workload costs, and checks scaling constraints defined in policies.

Points to consider:

– Intersight Workload Optimizer can generate scaling actions for the following VMs, but cannot execute the actions automatically:

• VMs with local SSDs

Intersight Workload Optimizer 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.

– Intersight Workload Optimizer can recommend scaling VMs to instance types with existing CUD capacity. CUD purchase recommendations will be introduced in a future release.

NOTE:
Intersight Workload Optimizer 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, Intersight Workload Optimizer can still recommend actions to apply CUD to other projects in your environment.

– Since all GCP compute tiers have the same net throughput capacity, Intersight Workload Optimizer will not generate scaling actions in response to net throughput.

– Intersight Workload Optimizer does not recommend scaling actions for:

• Spot VMs

NOTE:
Intersight Workload Optimizer 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, Intersight Workload Optimizer 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, Intersight Workload Optimizer will recommend that you reconfigure the VM.

- **Delete Volume**

  Delete unattached GCP volumes as a cost-saving measure. Intersight Workload Optimizer generates an action immediately after discovering an unattached volume.

  Intersight Workload Optimizer 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, Intersight Workload Optimizer discovers these volumes, but does not show actions or costs in charts.
  - Intersight Workload Optimizer 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 the subscription and Intersight Workload Optimizer discovers the resources available to you through that service principal.

Through Azure service principal targets, Intersight Workload Optimizer automatically discovers the subscriptions to which the service principal has been granted access in the Azure portal. This in turn creates a derived target for each subscription that inherits the authorization provided by the service principal (e.g. contributor). You cannot directly modify a derived target, but Intersight Workload Optimizer validates the target and discovers its inventory as it does with any other target.

NOTE:
In addition to service principal targets, you must also specify billing targets so Intersight Workload Optimizer can use custom pricing and discover reservations. When you configure billing targets, Intersight Workload Optimizer uses that richer pricing information to calculate workload size and reservation coverage for your Azure environment.

Requirements

To claim an Azure Service Principal target, you must meet the following requirements:

- Set up your Azure Service Principal subscription to grant Intersight Workload Optimizer the access it needs.

  To set up the Azure subscription, you must access the Administrator or Co-Administrator Azure Portal (`portal.azure.com`). Note, *this access is only required for the initial setup*. Intersight Workload Optimizer does not require this access for regular operation. See Azure Subscription Setup *(on page 27)*.
- Claim the target with the credentials that result from the subscription setup (Tenant ID, Client ID, etc.).
- Azure Resource Manager

  Intersight Workload Optimizer requires the Azure Resource Manager deployment and management service. This provides the management layer that Intersight Workload Optimizer uses to discover and manage entities in your Azure environment.

NOTE:
Intersight Workload Optimizer does not discover Azure Classic virtual machines, because they do not utilize the Azure Resource Manager.

After you claim an Azure target, you must allow at least 30 minutes of discovery time for Intersight Workload Optimizer to fully load Resource Group information.

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:

- Intersight Workload Optimizer 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*, Intersight Workload Optimizer discovers the related app instances. In the supply chain, app instances appear as 'App Component Specs'. Intersight Workload Optimizer generates actions to scale these plans to optimize app performance.
- For plans associated with the other types of apps, Intersight Workload Optimizer does not generate scale actions or discover the related app instances.
- For plans that are not associated with any type of app, Intersight Workload Optimizer 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 Required Permissions *(on page 27)*.

Supported Regions – Azure

Intersight Workload Optimizer 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	

Region Code	Region Name	Notes
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	
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/

Region Code	Region Name	Notes
		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)

Azure Subscription Setup

To claim an Azure Service Principal target, you must first set up your Azure subscription to grant Intersight Workload Optimizer the access it needs. This includes:

- Setting up the Azure Required Permissions *(on page 27)*
- Registering Intersight Workload Optimizer with Azure Active Directory *(on page 28)*

  This gives you the Application ID and Tenant ID.
- Creating the Client Secret Key and Permissions *(on page 28)*
- Enabling Intersight Workload Optimizer Access to Subscriptions *(on page 29)*

Azure Required Permissions

Intersight Workload Optimizer interacts with Azure targets through an Azure AD Application/Service Principal. As part of subscription setup, you must assign a permission level through one of the built-in Azure roles. Please review the information below for details before you proceed:

- Action Execution Access

  For action execution access, the user must have the `Owner` or `Contributor` role. `Contributor` is the least privileged role that enables Intersight Workload Optimizer to take actions on your Azure environment, including manually or automatically scaling VMs across instance types or automating VM stop and start.
- Read-Only Access

  Use the combined `Reader` plus `Storage Account Contributor` for the minimum combination required for Intersight Workload Optimizer to discover and access metrics across your Azure environment. `Storage Account Contributor` is required to access the Storage Account keys and establish a connection that can retrieve VM memory statistics.

  You can also use a combination of `Reader` on the subscription, and `Storage Account - List Keys` on the storage account where memory metrics are stored. To set this up, you must create the `Storage Accout - List Keys` role via the Azure CLI or APIs.

  For example, you can edit the following to add your keys:

```
{
  "Name":"Turbonomic Storage Key Access",
  "IsCustom": true,
  "Description": "Can list storageAccount keys.",
  "Actions":[
    "Microsoft.Storage/storageAccounts/listkeys/action"
  ],
  "NotActions": [
  ],
  "DataActions": [
  ],
  "NotDataActions": [
```

```
  ],
  "AssignableScopes": [
    "/subscriptions/<INSERT_SUBSCRIPTION_ID>"
  ]
}
```

If you save this JSON to a file named `listkeys.json`, you can execute the following command to create the permission:

```
az role definition create --role-definition listkeys.json
```

For other approaches to specify enhanced security, please contact your support representative.

Registering Intersight Workload Optimizer with Azure Active Directory

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. Intersight Workload Optimizer connects to an Azure target via an App registration.

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

To register the application:

1. Log into the Microsoft Azure Portal.

   You can log in at https://portal.azure.com.
2. Open the blade to register the Intersight Workload Optimizer application.
   - Navigate to **Azure Active Directory**
   - Under **Manage**, click **App registrations**
   - Click the **New registration** button

   This opens a new blade for registering the application.
3. Register Intersight Workload Optimizer as an application in the new blade.

   Enter the required details in the blade:
   - Name

     This can be any name that you want, e.g. `IWO`.
   - Supported Account Types

     Be sure to select the default option (Default Directory).

   Then click **Register**.
4. Note the details of the registered application.

   You have now created the App Registration for Intersight Workload Optimizer. Make a note of the **Application (client) ID** and **Directory (Tenant) ID** values. You will need these values later when you claim the Azure target in Intersight Workload Optimizer.

Creating the Client Secret Key and Permissions

To create the secret key:

1. Navigate to the **Certificates and Secrets** section of your registered app.
2. Create the client secret.

   Click **+ New client secret** to start. In the **Add a client secret** area, provide the required information:
   - **Description:** Give `IWO` or any other name.
   - **Expires:** Choose `Never`

   To create the secret, click **Add**.

When you click **Add**, the blade generates a secret for you to use when claiming the target in Intersight Workload Optimizer.

IMPORTANT:
Make sure to copy the secret. It **will not** be displayed again after you leave this page.

To set the required permissions:

1. Navigate to the **API Permissions** section of your registered application.
2. Click **+Add a permission** and select **Azure Service Management**.
3. Select **Delegated permissions** and check the box for **use_impersonation**.
4. When you have made your settings, click **add permissions**

   .

Enabling Intersight Workload Optimizer Access to Subscriptions

The final step in setting up your Azure subscription is to add the Active Directory Application (that you created above) to each and every subscription that you want Intersight Workload Optimizer to manage. Note that if you need to, you can assign different permissions to each subscription (for example, Service Principal).

To add the application to your subscriptions, first display the subscriptions that are set up for your tenant ID. In the Azure Portal, search for the **Subscriptions** button, and click it. This should display all the subscriptions that are under your Tenant (Directory) ID.

NOTE:
If you only see a single subscription, be sure to uncheck the option, **Show only subscriptions selected in the global subscription filter**.

Then for each subscription within the tenant:

1. Add a role assignment to the subscription.

   Click **Control Access (IAM)**. Then click **Add** (at the top of the control area) and select **Add role assignment**.
2. Select the role to add to this subscription.

   Select the role as one of:

   - `Owner`
   - `Contributor`
   - A combined role of `Reader` plus `Storage Account Contributor`

   For more information about permissions, see the explanation in Azure Required Permissions *(on page 27)*.
3. Map the application to this permission.

   In the **Select** field, give the name that you used when you registered Intersight Workload Optimizer with Azure Active Directory (see Registering Intersight Workload Optimizer with Azure Active Directory *(on page 28)*).

   When you have entered the application name, click **Save**.

   After you assign the application and role to the subscription, you should see the application as one of the subscription users in the list of subscription roles. You should also see an **Added User** notification that says the application was added as a user to the current subscription.
4. Repeat these steps for every subscription that you want Intersight Workload Optimizer to discover and manage through this Service Principal subscription.

Claiming Azure Targets

To claim an Azure target, you should have the following information at hand:

- Tenant/Directory ID

  You should have noted the Tenant ID when you registered the Intersight Workload Optimizer application. For more information, see Registering Intersight Workload Optimizer with Azure Active Directory *(on page 28)*.
- Application (Client) ID

  You should have noted the Client ID when you registered the Intersight Workload Optimizer application. For more information, see Registering Intersight Workload Optimizer with Azure Active Directory *(on page 28)*.
- Client Secret Key

  You should have saved the secret key when you created the secret key for the Intersight Workload Optimizer application. For more information, see Creating the Client Secret Key and Permissions *(on page 28)*.

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

- Name

  The display name that will be used to identify the target in the Intersight Workload Optimizer Targets List. This is for display in the user interface, only. The name you provide here does not need to match any internal name.
- Tenant ID

  The ID of the tenant that contains subscriptions to be managed by Intersight Workload Optimizer.
- Client App ID

  The Client/App ID of the App Registration that gives Intersight Workload Optimizer access to resources in your Azure subscription.
- Client Secret Key

  The secret key for the App Registration.

After you successfully claim the Azure target, you should see two types of entries for Azure in the Targets list:

- Azure Service Principal Target

  This target includes the IDs provided earlier and it will list all the subscriptions the service principal was assigned to. You can edit this target if necessary.
- Azure Target

  This entry represents the individual subscription discovered by the Azure Service Principal Target. You cannot edit this type of target entry.

Azure Discovery and Management

After you have set up the Azure Service Principal subscription and then claimed it as a target, Intersight Workload Optimizer can discover the Azure environment.

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

Azure Classic Virtual Machines

Intersight Workload Optimizer requires the Azure Resource Manager deployment and management service. This provides the management layer that Intersight Workload Optimizer uses to discover and manage entities in your Azure environment.

Because they do not utilize the Azure Resource Manager, Intersight Workload Optimizer does not discover Azure Classic virtual machines.

Locked Storages and Resource Groups

NOTE:
After you claim an Azure target, you must allow at least 30 minutes of discovery time for Intersight Workload Optimizer to fully load Resource Group information.

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

- Locked resource group

  Intersight Workload Optimizer 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

  Intersight Workload Optimizer discovers all the entities in the resource group except the locked storage. It also discovers the resource group.

Accessing Reservations

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

Cloud Instance Family Support

In the user interface, you can see the instance types that Intersight Workload Optimizer supports.

1. Navigate to **More > 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.

Intersight Workload Optimizer 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 Azure environment. With memory metrics, Intersight Workload Optimizer can generate actions that not only boost performance but also maximize your savings.

For Intersight Workload Optimizer to collect memory statistics in Azure, you must enable the collection of these statistics on the VMs in your environment. You can do this as you deploy your VMs, or you can enable the counters on VMs you have already deployed. For more information, see Enabling Collection of Memory Statistics: Azure *(on page 148)*.

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the cloud target supply chain:

Entity Type	Commodity
Virtual Machine (Azure)	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) ■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS ■ I/O Throughput The throughput to the underlying storage for the entity

Entity Type	Commodity
	Measured in KB/s ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Virtual Machine Spec (Azure App Service Plan)	■ Virtual Memory (VMem) The utilization of VMem allocated to an App Service plan ■ Virtual CPU (VCPU) The utilization of VCPU allocated to an App Service plan ■ Storage Amount The amount of Azure storage utilized by an App Service plan ■ Number of Replicas The total number of VM instances underlying an App Service plan
App Component Spec (Azure App Service)	■ Response Time The elapsed time between a request to an App Component Spec and the response to that request ■ Virtual CPU (VCPU) The percentage of total CPU time utilized by a given app
Database Server	■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)
Database	**NOTE:** The resources that Intersight Workload Optimizer can monitor depend on the pricing model in place for the given database entity. ■ DTU Pricing Model – DTU DTU capacity for the database. DTU represents CPU, memory, and IOPS/IO Throughput bundled as a single commodity. – Storage Storage capacity for the database. ■ vCore Pricing Model – Virtual Memory (VMem) The utilization of VMem allocated to the Database instance – Virtual CPU (VCPU) The utilization of VCPU allocated to the Database instance – Storage Access (IOPS) The rate of input and output operations per second utilized by the Database instance – Throughput The throughput utilization of transaction log write IO available to the Database instance – Storage Storage capacity for the database.

Entity Type	Commodity
	Intersight Workload Optimizer drives scaling actions based on the utilization of these resources, and treats the following limits as constraints when it makes scaling decisions: ■ Maximum concurrent sessions Maximum number of database connections at a time. ■ Maximum concurrent workers Maximum number of database processes that can handle queries at a time.
Volume	■ Storage Access The percentage of the volume's capacity for storage access operations (measured in IOPS) that is in use. ■ IO Throughput The percentage of the volume's capacity for IO throughput (measured in MB/s) that is in use. ■ IO Throughput Read The percentage of the volume's capacity for IO throughput Read (measured in MB/s) that is in use. ■ IO Throughput Write The percentage of the volume's capacity for IO throughput Write (measured in MB/s) that is in use.

Actions

Intersight Workload Optimizer 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 Intersight Workload Optimizer 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

  Intersight Workload Optimizer 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*.

Microsoft Azure Billing Targets

The Intersight Workload Optimizer Azure Billing target discovers your Azure billing account and related subscriptions. The target can access your billing data through your Enterprise Agreement (EA) offer ID.

NOTE:
Use Microsoft Azure Billing targets for non-government accounts.

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 Microsoft Azure Billing target, the Top Billed Cost by Account, Top Billed Cost by Service, and Top Billed Cost by Service Provider widgets do not display Azure billing data. This will be supported in a future release.

The target discovers:

- Billing Organization

  Intersight Workload Optimizer discovers the billing account and related subscriptions associated with your EA offer ID.
- Azure Reservations

  Intersight Workload Optimizer discovers all reservations that are charged under your billing account.
- Billing Costs

  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.

Prerequisites

You must have set up a daily Cost Export.

Billing targets use Cost and Usage reports. In order for Intersight Workload Optimizer 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 set up.

- 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`

Claiming an Azure Billing Target

NOTE:
You should remove any corresponding EA targets prior to adding an Azure Billing target.

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

  - EA Billing Account

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

    `12345678`

- Cost Export Name

  The name of the cost export.

  To find the cost export name, log into the Azure portal and navigate to **Cost management + Billing**. From there, select **Exports** to display a list of the cost exports you have created. For more information about setting up cost exports, see Prerequisites *(on page 34)*.

- Directory (Tenant) ID

  The ID of the tenant that contains subscriptions to be managed with Intersight Workload Optimizer. This should match the Directory ID that you give for the associated Service Principal Azure target.

- 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.

- Application (Client) ID

  The Client/App ID of the App Registration that gives Intersight Workload Optimizer access to resources in your Azure subscription. This should match the Client/App ID that you give for the associated Service Principal Azure 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 Azure target.

Note that it takes 24 hours after the target has been successfully added for any RI or Discount Coverage information to display in the UI.

Azure Billing Target Permissions

To configure Azure Billing targets, you create service principal accounts with roles that grant Intersight Workload Optimizer 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.

Intersight Workload Optimizer Functionality	Required Permissions for Service Principal
Monitoring	`Enrollment Reader` role (EA) ■ 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.

Firewall Access

If you run Intersight Workload Optimizer behind a firewall, you must configure Intersight Workload Optimizer 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`

Actions

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

Monitored Resources

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

Microsoft Enterprise Agreement

You can configure Intersight Workload Optimizer to manage Azure subscriptions within the context of an Enterprise Agreement (EA). An EA target enables Intersight Workload Optimizer to use custom pricing and discover reservations. When you configure an EA target, Intersight Workload Optimizer uses that richer pricing information to calculate workload size and reservation coverage for your Azure environment.

NOTE:
Use Microsoft Enterprise Agreement for government accounts.

To enable Intersight Workload Optimizer 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 24)*.

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, Intersight Workload Optimizer 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.

Claiming 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 Enterprise Agreement enrollment number (found in your EA admin account at ea.azure.com)
- API Key

  The API Access Key for the Enterprise Agreement (found in your EA admin account at ea.azure.com)

When you add the target and it validates, Intersight Workload Optimizer:

- 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: Intersight Workload Optimizer 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 24)*.

Azure Enterprise Agreements

You can configure Intersight Workload Optimizer 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, Intersight Workload Optimizer uses that richer pricing information to calculate workload placement and reservations coverage for your Azure environment.

To enable Intersight Workload Optimizer 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 24)*.

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 Intersight Workload Optimizer 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, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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. Intersight Workload Optimizer combines this information to track:

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

NOTE:
This release of Intersight Workload Optimizer 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, Intersight Workload Optimizer 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.

  Intersight Workload Optimizer 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 Intersight Workload Optimizer 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, Intersight Workload Optimizer recognizes Azure Hybrid Benefit (AHUB) savings for reservations and on-demand workloads. The costs you see in Intersight Workload Optimizer 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).

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, Intersight Workload Optimizer targets Kubernetes clusters. Intersight Workload Optimizer supports target clusters managed on Kubernetes v1.8 or higher, whether the clusters are managed directly via kubeadm, or via other platforms including:

- OpenShift
- Pivotal Kubernetes Service
- Amazon Elastic Kubernetes Service (EKS)
- Azure Kubernetes Service (AKS)
- Google Kubernetes Engine (GKE)

With Cloud Native targets, Intersight Workload Optimizer 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 Intersight Workload Optimizer:

- Cisco AppDynamics *(on page 82)*
- Dynatrace *(on page 91)*
- New Relic *(on page 87)*

Claiming a Kubernetes Platform Target

To claim this target for a Kubernetes cluster, you first install the Intersight Workload Optimizer Kubernetes Collector on your Kubernetes cluster. The installation process generates a Device ID and a Claim Code that you can then use to add the collector as a target to manage the Kubernetes cluster.

The Intersight Workload Optimizer platform gathers information from your Kubernetes or OpenShift environment via the collector that you install on your Kubernetes cluster. The collector collects information from your environment and passes it to Intersight Workload Optimizer. As it generates actions, Intersight Workload Optimizer then uses the collector to execute those actions in your cluster. In this way, Intersight Workload Optimizer users can execute actions from the user interface, policies can set up actions to execute automatically, and Intersight Workload Optimizer can automatically execute groups of related actions on the workloads in a container spec.

NOTE:
You must install a different collector on each Kubernetes cluster you want Intersight Workload Optimizer to manage.

You install the collector via a Helm chart. For installation instructions, see Installing the Intersight Workload Optimizer Kubernetes Collector *(on page 45)*. The last installation step is to register the collector. This generates a Device ID and a Claim Code that you can then use to add the collector as a target. As you install the collector, you should record these values.

To claim a Kubernetes target, select **Cloud Native > Kubernetes** on the Target Configuration page and provide the following information:

- Device ID

  This identifies the Kubernetes Collector you have installed for the given cluster. When you register the collector, this is returned as the `SerialNumber` token.
- Claim Code

  This authorizes the connection between your Intersight Workload Optimizer account and the collector. When you register the collector, this is returned as the `SecurityToken`.

Actions

Intersight Workload Optimizer recommends actions for the Kubernetes container platform supply chain as follows.

Entity	Action
Service	**None** No actions are recommended at this level of the supply chain. Instead, actions that affect the service are generated and executed on underlying entities.
Application Component	**Suspend** Application components are suspended due to a node (virtual machine) suspension **APM Actions** Application components may also receive other actions as part of APM integration related to those use cases. For example, a `Resize Heap` action from an underlying AppDynamics integration. See the Target Configuration Guide documentation for the appropriate technology to discover what actions may be available.
Container	**Resize Container Up/Down** With `Merged Actions` enabled, individual Container actions will be recommend only and the resize will be reflected as an action on the Workload Controller entity. **Suspend** Containers are suspended due to a node (virtual machine) suspension
Container Pod	**Move Pod** Pods will be moved across nodes (Virtual Machines). **Suspend** Container Pods are suspended due to a node (virtual machine) suspension
Container Spec	**None** No actions are recommended at this level of the supply chain. This entity maintains the history of all replicas, or instances of pods for this container specification.
Workload Controller	**Resize Container** With `Merged Actions` enabled, this is a single resize action representing all resize actions for containers associated to a specific workload controller.

Entity	Action
Namespace	**None** No actions are recommended at this level of the supply chain. Namespace Quotas are constraints to container resizing actions.
Virtual Machine (Node)	**Provision Additional Resources** The following resources may be provisioned: ■ VMem ■ VCPU ■ VMem Requests ■ VCPU Requests ■ Number of Consumers **Suspend** Nodes (virtual machines) may be suspended. **Infrastructure-dependent Actions** Depending on the technology the node is stitched to for underlying infrastructure, there may be additional actions: On-prem VMware: ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine
Volume	**None** No actions are recommended at this level of the supply chain. These entities will be stitched to public cloud storage volumes.

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the Kubernetes container platform supply chain:

Entity	Commodity
Service	**Response Time** Response time of the service, measured in ms. This commodity is populated via APM or DIF integrations. **Transactions** Transaction utilization, measured in transactions per second. This commodity is populated via APM or DIF integrations.
Application Component	**Various Commodities** The commodities monitored and the values received for those commodities at the application component level is dependent on the APM integration used. See the Target Configuration Guide documentation for the appropriate technology to discover what data will be reported.
Container	**VMem** The virtual memory utilized by the container against the memory limit (if no limit is set, then node capacity is used). Measured in Megabytes (MB) **VMem Request** If applicable, the virtual memory utilized by the container against the memory request. Measured in Megabytes (MB)

Entity	Commodity
	VCPU The virtual CPU utilized by the container against the CPU limit (if no limit is set, then node capacity is used). Measured in millicores (mCores) **VCPU Request** If applicable, the virtual CPU utilized by the container against the CPU request. Measured in millicores (mCores) **VCPU Throttling** 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, *used* and *utilization* values reflect the actual throttling percentage, while *capacity* value is always 100%.
Container Pod	**VMem** The virtual memory utilized by the pod against the node physical capacity. Measured in Megabytes (MB) **VCPU** The virtual CPU utilized by the pod against the node physical capacity. Measured in millicores (mCores) **VMem Request** The virtual memory request allocated by the pod against the node allocatable capacity. Measured in Megabytes (MB) **VCPU Request** The virtual CPU request allocated by the pod against the node allocatable capacity. Measured in millicores (mCores) **VMem Request Quota** If applicable, The amount of virtual memory request the pod has allocated against the namespace quota. Measured in Megabytes (MB) **VCPU Request Quota** If applicable, The amount of virtual CPU request the pod has allocated against the namespace quota. Measured in millicores (mCores) **VMem Limit Quota** If applicable, The amount of virtual memory limit the pod has allocated against the namespace quota. Measured in Megabytes (MB) **VCPU Limit Quota** If applicable, The amount of virtual CPU limit the pod has allocated against the namespace quota. Measured in millicores (mCores)
Container Spec	**VMem** The virtual memory historically utilized by any containers run for this workload against the memory limit (if no limit is set, then node capacity is used). Measured in Megabytes (MB) **VCPU** The virtual CPU historically utilized by any containers run for this workload against the CPU limit (if no limit is set, then node capacity is used). Measured in millicores (mCores) **VMem Request**

Entity	Commodity
	If applicable, the virtual memory historically utilized by any containers run for this workload against the memory request. Measured in Megabytes (MB) **VCPU Request** If applicable, the virtual CPU historically utilized by any containers run for this workload against the CPU request. Measured in millicores (mCores)
Workload Controller	**VMem Request Quota** If applicable, The amount of virtual memory request the pod has historically allocated for this workload against the namespace quota. Measured in Megabytes (MB) **VCPU Request Quota** If applicable, The amount of virtual CPU request the pod has historically allocated for this workload against the namespace quota. Measured in millicores (mCores) **VMem Limit Quota** If applicable, The amount of virtual memory limit the pod has historically allocated for this workload against the namespace quota. Measured in Megabytes (MB) **VCPU Limit Quota** If applicable, The amount of virtual CPU limit the pod has historically allocated for this workload against the namespace quota. Measured in millicores (mCores)
Namespace	**VMem Request Quota** The total amount of virtual memory request for all pods allocated to the namespace against the namespace quota. Measured in Megabytes (MB) **VCPU Request Quota** The total amount of virtual CPU request for all pods allocated to the namespace against the namespace quota. Measured in millicores (mCores) **VMem Limit Quota** The total amount of virtual memory limit for all pods allocated to the namespace against the namespace quota. Measured in Megabytes (MB) **VCPU Limit Quota** The total amount of virtual CPU limit for all pods allocated to the namespace against the namespace quota. Measured in millicores (mCores)
Virtual Machine (Node)	**VMem** The virtual memory utilized by the node against the memory allocated to the hosting virtual machine. Measured in Megabytes (MB) **VCPU** The virtual CPU utilized by the node against the CPU allocated to the hosting virtual machine. Measured in Megahertz (Mhz) **VMem Request** The total amount of virtual memory allocated to pods with memory request against the allocatable capacity of the node. Measured in Megabytes (MB) **VCPU Request**

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

Installing the Intersight Workload Optimizer Kubernetes Collector

To install the Intersight Workload Optimizer Kubernetes Collector, you deploy it on a node in your Kubernetes cluster. From that node, the collector uses `kubelet` to reach all the other pods in the cluster. To download the latest version of the Intersight Kubernetes Collector, go to Cisco Software Download.

The collector installs as two identical pods. This supports High Availability (HA), where one pod can take over if the currently active collector pod crashes.

You deploy the collector on one node per cluster, or one collector per control plane when using stretch clusters. The collector runs with a service account that has the `cluster-admin` role. This role enables the collector to execute Intersight Workload Optimizer actions within your Kubernetes cluster.

To communicate with Intersight Workload Optimizer, the collector installs with its own Device Connector. The device connector provides a secure way for the collector to send information and receive control instructions from the Cisco Intersight portal, using a secure Internet connection.

Installation Requirements

To use the Intersight Workload Optimizer Kubernetes Collector, your environment must meet the following requirements:

- Kubernetes or OpenShift version:
  - Kubernetes version 1.8 or higher
  - OpenShift release 3.4 or higher
- Helm v2 or v3 installed:

  To deploy the collector, you will use Helm to install a chart in the Kubernetes cluster.

  The installation instructions assume you have Helm v2 or v3 installed and configured to install the chart on the node where you want the collector to run. For more information about the Help client, see HELM.

  For Helm v2, you must also have Tiller installed. Tiller requires the `cluster-admin` role to install and run collector charts, and needs to run with a service account with `ClusterRole` access. For details about role-based access, see:
  - https://helm.sh/docs/topics/rbac/
  - https://github.com/fnproject/fn-helm/issues/21
- Network requirements:

  The collector pods must have access to the kubelet on every node in the cluster. This access can be via one of:
  - `https + port=10250 (default)`

  - `http + port=10255`

- Device Connector port requirements:

  The device connector provides a secure way for the collector to send information and receive control instructions from the Cisco Intersight portal. The following table lists the port numbers that must be open for device connector communication:

Port:	Protocol:	Description:
`443`	TCP/UDP	Required for communication between: – The device connector and the user's Web browser. – The device connector and the Kubernetes endpoints.
`80`	TCP	This port is optional for normal operation, but is required for initial monitoring of the device connector setup and when using the one-time device connector upgrade. This port is not used if the device connector is at the minimum supported version.

- Compute and storage requirements:

  The collector pod typically runs well with no more than:

  - 512 Mg Memory
  - One core or 1 GHz CPU
  - 10 GB of volume space

Deploying the Intersight Workload Optimizer Kubernetes Collector

To deploy the collector:

1. Download the Helm chart to your node cluster.

   To get the Helm Chart files, download them from Cisco Software Download.

2. Create a namespace for the collector.

   If you are installing many instances of the collector on many node clusters, it can be convenient to use the same namespace for each deployment. Execute the following command, where **`iwo-collector`** is the namespace name for this example (you can use any valid namespace name):

   ```
   kubectl create namespace iwo-collector
   ```

3. Execute the install command.

   When you execute the command, you will specify:

   - `name`: The release name of the installed pod.

     The name must not exceed 20 characters in length.
   - `namespace`: The namespace the pod is installed under.
   - `iwoServerVersion`: The main version family for the Intersight Workload Optimizer that will manage your cluster. For this installation, give the value, `8.7`.

     As we add features to Intersight Workload Optimizer, you might be instructed to update this value to a higher version.
   - `collectorImage.tag`: The version of the collector that you are installing. For this installation, give the value, `8.7.5.1`.

     As we add features to the collector, you might be instructed to update this value to a higher version.
   - `targetName`: A name that identifies the cluster you are installing onto. This can be any name.

The use interface uses this name as it displays each managed cluster in the list of targets. It also uses this name in other places where it displays data about the cluster.

The following sample commands assume these parameter values:

- `name = my-iwo-k8s-collector`
- `namespace = iwo-collector`
- `iwoServerVersion = 8.7`
- `collectorImage.tag = 8.7.5.1`
- `targetName = my-k8s-cluster`

Before you actually install the collector, you should use the dry-run feature. Execute the command:

- Helm v2:

```
helm install --dry-run --debug <Chart_Location> --name my-iwo-k8s-collector --namespace iwo-collector --set iwoServerVersion=8.7 --set collectorImage.tag=8.7.5.1 --set targetName=my-k8s-cluster
```

- Helm v3:

```
helm install --dry-run --debug my-iwo-k8s-collector <Chart_Location> --namespace iwo-collector --set iwoServerVersion=8.7 --set collectorImage.tag=8.7.5.1 --set targetName=my-k8s-cluster
```

Inspect the output to make sure the results are as you expect. If the output is correct, execute the installation:

- Helm v2:

```
helm install <Chart_Location> --name my-iwo-k8s-collector --namespace iwo-collector --set iwoServerVersion=8.7 --set collectorImage.tag=8.7.5.1 --set targetName=my-k8s-cluster
```

- Helm v3:

```
helm install my-iwo-k8s-collector <Chart_Location> --namespace iwo-collector --set iwoServerVersion=8.7 --set collectorImage.tag=8.7.5.1 --set targetName=my-k8s-cluster
```

Note that you can provide custom values for the installation. The parameters you can access and optionally change are:

Parameter	Default Value	Required / Opt to Change	Parameter Type
connectorImage.repository	intersight/pasadena	optional	path to connector repo
connectorImage.tag	1.0.9-24	optional	connector image tag
connectorImage.pullPolicy	IfNotPresent	optional	
collectorImage.repository	turbonomic/kubeturbo	optional	path to repo
collectorImage.tag	8.7.5.1	optional	IWO Kubernetes Probe image tag
collectorImage.pullPolicy	IfNotPresent	optional	
iwoServerVersion	8.7	required	number x.y
targetName	"Your_k8s_cluster"	optional Note that this is required for multiple clusters.	String The name you want to use to identify your cluster

Parameter	Default Value	Required / Opt to Change	Parameter Type
args.failVolumePodMoves	true	optional Change to false if you want to move pods which have volumes attached. The pod(s) will be down during the move.	boolean
args.kubelethttps	true	optional Change to false if k8s 1.10 or older.	boolean
args.kubeletport	10250	optional Change to 10255 if k8s 1.10 or older.	number
args.logginglevel	2	optional	number
args.stitchuuid	true	optional Change to false if IaaS is VMM, Hyper-V.	boolean
HANodeDetectors.nodeRoles	"\"master\""	Optional Used to automate policies to keep nodes of same role limited to 1 instance per ESX host or AZ.	string Uses regex. Values in quotes, and comma separated; e.g. `"master"` (default),`"worker"`,`"app"`, etc.

4. Verify the installation.

   After you execute the installation, you should give Helm enough time to install the collector. It installs two collector pods to support HA for the collector. Each pod contains two containers; one for the device connector, and one for the collector.

   After enough time has passed to start up the pods, you should verify that they are running. Execute the following command, where **`iwo-collector`** is the namespace for this collector:

   ```
   kubectl get pods -n iwo-collector
   ```

   The entry for the pod should have the following values:

   - NAME: `iwok8scollector-`**`my-iwo-k8s-collector`**`<Pod_ID>`, where **`my-iwo-k8s-collector`** is the name you provided when you installed the collector, and `<Pod_ID>` is a generated ID value.
   - READY: `2/2`
   - STATUS: `Running`

   For example, the output should be similar to:

   ```
   NAME                                                  READY   STATUS    RESTARTS   AGE
   iwok8scollector-my-iwo-k8s-collector-57fcb8b874-s5ch8   2/2     Running   0          12s
   iwok8scollector-my-iwo-k8s-collector-57fcb8b874-c7dd2   2/2     Running   0          12s
   ```

   You will need one of the full pod names for the next step. Either pod will do.

5. Register the collector to get its Device ID and Claim Code.

   When you register the collector, you will forward its port 9110, and then you will connect to your Intersight instance to get the registration tokens. If you need a proxy for a connection external to your cluster or network, then you can run a command to enable a proxy in the collector.

   To register the collector:

   - Forward the pod's port 9110:

Execute the command, where **`iwo-collector`** is the namespace you defined for this collector, and **`my-iwo-k8s-collector-57fcb8b874-s5ch8`** is the full pod name:

```
kubectl -n iwo-collector port-forward my-iwo-k8s-collector-57fcb8b874-s5ch8 9110
```

- (Optional) Configure the proxy connection from the collector to `intersight.com`.

  If your proxy does not require authentication, execute the following command:

  ```
  curl -XPUT http://localhost:9110/HttpProxies -d '{"ProxyType":"Manual", "ProxyHost":"My_Proxy_Server", "ProxyPort":"My_Proxy_Port"}'
  ```

  If your proxy requires authentication using usernname/password credentials, execute the following command:

  ```
  curl -XPUT http://localhost:9110/HttpProxies -d '{"ProxyType":"Manual", "ProxyUsername":"<username>", "ProxyPassword":"<password>", "ProxyHost":"My_Proxy_Server", "ProxyPort":"My_Proxy_Port"}'
  ```

  Where:

  - **`My_Proxy_Server`** is the address of your proxy server

    Note that you must *not* include the HTTP protocol in the proxy address. For example, if your proxy is located at `https://proxy-was.esl.cisco.com`, you would give the following address:

    ```
    proxy-was.esl.cisco.com
    ```

  - **`My_Proxy_Port`** is the port your proxy server uses

- Get the Device ID and Claim Code:

  Execute the following commands:

  - Get the Device ID:

    ```
    curl -s http://localhost:9110/DeviceIdentifiers
    ```

    The command output should be similar to the following, where `"ID":` is the Device ID value:

    ```
    [
      {
        "Id": "22284c13-xxxx-yyyy-zzzz-93a14e4de07f"
      }
    ]* Closing connection 0
    ```

  - Get the Claim Code:

    ```
    curl -s http://localhost:9110/SecurityTokens
    ```

    The command output should be similar to the following, where `"Token":` is the Claim Code value:

    ```
    [
      {
        "Token": "26AEAECCDD67",
        "Duration": 599
      }
    ]* Closing connection 0
    ```

  Record these values. You will provide them as credentials when you claim the collector as a Kubernetes target.

  If either command returns an error similar to the following, confirm that the cluster and the collector can connect with `intersight.com`. This could indicate that you need to configure a proxy connection (see above).

```
{
  "code":"InternalServerError",
  "message":"Internal error while fetching Claim Code",
  "messageId":"",
  "messageParams":null,
  "traceId":"DCxxxxxxxxxxxxxxxxfc9e4de952584049"
}
```

Updating the Intersight Workload Optimizer Kubernetes Collector

We continue to improve the management of resources in a Kubernetes cluster. To take advantage of these improvements, you should be sure to update the collector, so it can pass new types of data to Intersight Workload Optimizer and so it can execute any newly added commands.

When you update the collector, you will specify:

- `namespace`: The namespace the pod is installed under.
- `collectorImage.tag`: The version of the collector that you are updating to. For this installation, give the value, `8.7.5.1`.

The following sample commands assume these parameter values:

- **`<Chart_Location>`** for the location of the Helm chart (for Helm V2)
- **`my-iwo-k8s-collector`** for the pod name (for Helm V3)
- `namespace = iwo-collector`
- `collectorImage.tag = 8.7.5.1`

To update the collector, execute the Helm command:

- Helm v2:

```
helm upgrade <Chart_Location> --namespace iwo-collector --reuse-values --set collectorImag
e.tag=8.7.5.1
```

- Helm v3:

```
helm upgrade my-iwo-k8s-collector --namespace iwo-collector --reuse-values \
            --set collectorImage.tag=8.7.5.1
```

Removing the Intersight Workload Optimizer Kubernetes Collector

To remove the collector from your cluster, execute one of the following Helm commands, where **`iwo-collector`** is the namespace you used for the release, and **`my-iwo-k8s-collector`** is the release name:

- Helm v2:

```
helm uninstall -n iwo-collector my-iwo-k8s-collector
```

- Helm v3:

```
helm delete -n iwo-collector my-iwo-k8s-collector
```

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 Intersight Workload Optimizer will discover the managed application servers.

NOTE:
As it manages your applications environment, Intersight Workload Optimizer discovers connected application components to stitch them into a supply chain of entities. For connections that are made by name and not IP address, Intersight Workload Optimizer 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.

Apache Tomcat

Intersight Workload Optimizer supports connecting to individual Tomcat targets. Intersight Workload Optimizer 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.

  Intersight Workload Optimizer discovers Tomcat servers that are running on VMs or containers. The hosting VM or container must already be in your Intersight Workload Optimizer 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 40)*

- For information about hypervisor targets, see Hypervisor Targets *(on page 105)*

Configuring JMX Remote Access

Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer target.

  If you set the target scope, Intersight Workload Optimizer scans each VM within that group or cluster and tries to connect to the target over the specified port. Intersight Workload Optimizer 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, Intersight Workload Optimizer will require all database servers hosted on the VMs in the selected scope to be a valid target. If Intersight Workload Optimizer is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

Actions

Intersight Workload Optimizer recommends actions for the application supply chain as follows.

Entity Type	Action
Application Component (Tomcat Application)	■ Resize Heap Recommendation only. ■ Resize Thread Pool Recommendation only. ■ Resize Connection Capacity Recommendation only.
Virtual Machines	■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine ■ Suspend VM ■ Provision VM

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the application server supply chain:

Entity Type	Commodity
Application Component	■ Virtual Memory (VMem) The utilization of the VMem consumed from the hosting VM Measured in Kilobytes (KB)

Entity Type	Commodity
	■ Virtual CPU (VCPU) The utilization of VCPU consumed from the hosting VM Measured in Megahertz (MHz) ■ Transactions The utilization of the allocated transactions per second for the given application Measured in transactions per second ■ Heap The utilization of the application server's heap Measured in Kilobytes (KB) ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Threads The utilization of the server's thread capacity Measured in Threads ■ Connection The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ Remaining GC Capacity The percentage of server uptime spent not performing garbage collection Measured in uptime (%)
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

IBM WebSphere

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

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

You may need to make changes to your DNS or the file `/etc/resolv.conf` on the Intersight Workload Optimizer 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. Intersight Workload Optimizer 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

  Intersight Workload Optimizer discovers WebSphere servers that are running on VMs or containers. The hosting VM or container must already be in your Intersight Workload Optimizer 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 40)*
  - For information about container targets, see Hypervisor Targets *(on page 105)*

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 Intersight Workload Optimizer 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 Intersight Workload Optimizer target.

  If you set the target scope, Intersight Workload Optimizer scans each VM within that group or cluster and tries to connect to the target over the specified port. Intersight Workload Optimizer 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, Intersight Workload Optimizer will require all database servers hosted on the VMs in the selected scope to be a valid target. If Intersight Workload Optimizer is unable to authenticate a database server in the scope, the target will not validate and data will not be collected.

Actions

Intersight Workload Optimizer recommends actions for the application server supply chain as follows:

Entity Type	Action
Service	Intersight Workload Optimizer 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 Intersight Workload Optimizer can start a new application component to run another instance of the database application, and bind it to the service. On the other hand, if SQL requests drop off so that the load balancer only forwards requests to two of the databases, Intersight Workload Optimizer can suspend the dormant database and unbind it.
Application Component	■ Resize Heap This action can only be executed by Intersight Workload Optimizer when running in a domain controller. Standalone applications will see only a recommendation. ■ Resize Connection Capacity Recommendation only.
Virtual Machines	■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Reconfigure Storage ■ Suspend VM

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the application server supply chain:

Entity Type	Commodity
Service	■ Transactions The utilization of the allocated transactions per second for the given service Measured in transactions per second
Application Component	■ Virtual Memory (VMem) The utilization of the VMem consumed from the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of VCPU consumed from the hosting VM Measured in Megahertz (MHz) ■ Transactions The utilization of the allocated transactions per second for the given application Measured in transactions per second ■ Heap The utilization of the application server's heap Measured in Kilobytes (KB) ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Threads The utilization of the server's thread capacity Measured in Threads ■ Connection The utilization of the connection capacity. Only applicable to database servers

Entity Type	Commodity
	Measured in Connections ■ Remaining GC Capacity The percentage of server uptime spent not performing garbage collection Measured in uptime (%)
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

JVM Application

Intersight Workload Optimizer supports connecting to individual JVM Applications as targets. Intersight Workload Optimizer 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

  Intersight Workload Optimizer discovers JVM applications that are running on VMs or containers. The hosting VM or container must already be in your Intersight Workload Optimizer 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 40)*
  - For information about hypervisor targets, see Hypervisor Targets *(on page 105)*

Configuring JMX Remote Access

Intersight Workload Optimizer 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 Intersight Workload Optimizer target.

  If you set the target scope, Intersight Workload Optimizer scans each VM within that group or cluster and tries to connect to the target over the specified port. Intersight Workload Optimizer 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

Intersight Workload Optimizer recommends actions for the application supply chain as follows.

Entity Type	Action
Application Component (JVM Application)	■ Resize Heap Recommendation only.

Entity Type	Action
	■ 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	■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine ■ Suspend VM ■ Provision VM

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the application server supply chain:

Entity Type	Commodity
Application Component (JVM Application)	■ Heap The utilization of the application server's heap Measured in Kilobytes (KB) ■ Remaining GC Capacity The percentage of server uptime spent garbage collecting. Available when the JVM profiler is enabled. Measured in percentage of uptime (%)
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz)

SQL Server

NOTE:
This type of target can run as SaaS or in on-prem datacenters. When you claim the target, you can choose to turn ON or turn OFF **Connect through an Intersight Assist** as follows:

- If the target runs as SaaS:

  Turn OFF **Connect through an Intersight Assist**.

  You should be aware that for earlier versions of Intersight Workload Optimizer, to claim an AppDynamics target running as SaaS you were required to specify an Intersight Assist. If you claimed your target through an Assist, you can reclaim that target *without* using the Intersight Assist. To do that you must first delete the claimed target, and then claim the target anew with **Claim through an Intersight Assist** in the OFF position.
- If the target runs in an on-prem datacenter:

  Turn ON **Connect through an Intersight Assist**.

  To establish communication between this target on the datacenter and Intersight Workload Optimizer, you must:

  - Install an Intersight Assist appliance in the on-prem datacenter. The AppDynamics target must be accessible to the Intersight Assist appliance.
  - Connect the Intersight Assist instance with Cisco Intersight.
  - Log in to Cisco Intersight and claim the Intersight Assist instance as a target.
  - Claim the AppDynamics target with **Connect through an Intersight Assist** in the ON position.

    Intersight Assist provides a secure way for on-prem targets to send information to and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

Intersight Workload Optimizer 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 Intersight Workload Optimizer.

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 Intersight Workload Optimizer discovery
- To enable dynamic port discovery, the port used by the SQL Browser Service

Creating a Service User Account

The user account that Intersight Workload Optimizer 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 Intersight Workload Optimizer

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 Intersight Workload Optimizer UI
- Username

  Username for the account. This username must not include the AD domain
- AD Domain

  The Active Directory domain used by Intersight Workload Optimizer 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 Intersight Workload Optimizer target.

  If you set the target scope, Intersight Workload Optimizer scans each VM within that group or cluster and tries to connect to the target over the specified port. Intersight Workload Optimizer 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 port used to communicate with the browsing service. Intersight Workload Optimizer will obtain the SQLServer port for each instance running on each VM in the scope.
- SQLServer Port

  The SQL remote port. Intersight Workload Optimizer will use this port if there is no browsing service port specified, or if the browsing service is not available during discovery.

  NOTE:
  Intersight Workload Optimizer will connect to the port specified for the SQL browsing service first. If that connection fails, Intersight Workload Optimizer will connect using the SQLServer Port.
- Full Validation

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

Actions

Intersight Workload Optimizer recommends actions for the application supply chain as follows.

Entity Type	Action
Applications	Without separate targets to discover Guest OS Processes or Application Servers, Intersight Workload Optimizer 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, Intersight Workload Optimizer can also recommend provisioning a new host.
Virtual Machines	■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine ■ Suspend VM ■ Provision VM

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the application server supply chain:

Entity Type	Commodity
Database Server	■ DBMem The memory utilized by the database, as a percentage of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Intersight Workload Optimizer can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM. ■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ TransactionLog The utilization of the server's capacity for storage devoted to transaction logs Measured in Kilobytes (KB) ■ Cache Hit Rate The percentage of accesses that result in cache hits. Measured in a percentage of hits vs total attempts (%)
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

MySQL

NOTE:
This type of target can run as SaaS or in on-prem datacenters. When you claim the target, you can choose to turn ON or turn OFF **Connect through an Intersight Assist** as follows:

- If the target runs as SaaS:

  Turn OFF **Connect through an Intersight Assist**.

  You should be aware that for earlier versions of Intersight Workload Optimizer, to claim an AppDynamics target running as SaaS you were required to specify an Intersight Assist. If you claimed your target through an Assist, you can reclaim that target *without* using the Intersight Assist. To do that you must first delete the claimed target, and then claim the target anew with **Claim through an Intersight Assist** in the OFF position.
- If the target runs in an on-prem datacenter:

  Turn ON **Connect through an Intersight Assist**.

  To establish communication between this target on the datacenter and Intersight Workload Optimizer, you must:

  - Install an Intersight Assist appliance in the on-prem datacenter. The AppDynamics target must be accessible to the Intersight Assist appliance.
  - Connect the Intersight Assist instance with Cisco Intersight.
  - Log in to Cisco Intersight and claim the Intersight Assist instance as a target.
  - Claim the AppDynamics target with **Connect through an Intersight Assist** in the ON position.

    Intersight Assist provides a secure way for on-prem targets to send information to and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

To manage MySQL databases, Intersight Workload Optimizer 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 66)*

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 Intersight Workload Optimizer UI
- Username

  Username of the account Intersight Workload Optimizer uses to connect to the target.
- Password

  Password of the account Intersight Workload Optimizer 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 Intersight Workload Optimizer target.

  If you set the target scope, Intersight Workload Optimizer scans each VM within that group or cluster and tries to connect to the target over the specified port. Intersight Workload Optimizer 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, Intersight Workload Optimizer will use the MySQL default port of 3306.

- Full Validation

  When selected, Intersight Workload Optimizer will require all database servers hosted on the VMs in the selected scope to be a valid target. If Intersight Workload Optimizer 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	■ **Resize** (Recommend, only) – DBMem Resize actions are driven by metrics on the Database Server, instead of metrics on the hosting VM. Intersight Workload Optimizer uses DB memory and cache hit rate metrics to decide whether resize actions are necessary. – Connections
Virtual Machine	■ **Resize** – Resize resource capacity 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, Intersight Workload Optimizer verifies that the VM's cluster can adequately support the new size. If the cluster is highly utilized, Intersight Workload Optimizer will recommend a move action, taking into consideration the capacity of the new cluster and compliance with existing placement policies. For hypervisor targets, Intersight Workload Optimizer can resize vCPU by changing the VM's socket or cores per socket count. For details, see "VCPU Scaling Controls" in the *User Guide*. – Resize resource reservation 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 – Resize resource limit 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. ■ **Move** Move a VM due to: – High resource utilization on VM or host – Excess IOPS or latency in VStorage – Workload placement violation – Underutilized host (move VM before suspending host) ■ **Move VM Storage (Volume)**

Entity:	Actions:
	Move a VM's volume due to excess utilization of the current datastore, or for more efficient utilization of datastores in the environment. ■ **Reconfigure** Change a VM's configuration to comply with a policy. For hypervisor targets, Intersight Workload Optimizer can reconfigure VMs that violate vCPU scaling policies. For details, see "VCPU Scaling Controls" in the *User Guide*. ■ **Reconfigure VM Storage** Reconfigure overutilized storage resources by adding VStorage capacity. For underutilized storage resources, remove VStorage capacity.

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the application server supply chain:

Entity Type	Commodity
Database Server	■ DBMem The memory utilized by the database, as a percentage of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Intersight Workload Optimizer can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM. ■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ Cache Hit Rate The percentage of accesses that result in cache hits. Measured in a percentage of hits vs total attempts (%)
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM

Entity Type	Commodity
	Measured in milliseconds (ms)

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 fin 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer has an IP address of `10.10.123.45`. The following command grants privileges to that Intersight Workload Optimizer, 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 Intersight Workload Optimizer can use to log on
- Find the Service Name and port for the database

Prerequisites

- User permissions that grant access to Intersight Workload Optimizer through a specific user account. See Creating a Service User Account in Oracle *(on page 71)*.
- 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 Intersight Workload Optimizer target connection

Adding a Dynamic Performance View

In order to collect data from the Oracle database, Intersight Workload Optimizer 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 Intersight Workload Optimizer

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 Intersight Workload Optimizer user interface.
- Username/Password

  Valid client credentials for the database server. For Intersight Workload Optimizer 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 Intersight Workload Optimizer server.
- Scope

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

  If you set the target scope, Intersight Workload Optimizer scans each VM within that group or cluster and tries to connect to the target over the specified port. Intersight Workload Optimizer 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, Intersight Workload Optimizer will require all database servers hosted on the VMs in the selected scope to be a valid target. If Intersight Workload Optimizer 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	■ **Resize** (Recommend, only) – Connections – DBMem Resize actions are driven by metrics on the Database Server, instead of metrics on the hosting VM. Intersight Workload Optimizer uses DB memory and cache hit rate metrics to decide whether resize actions are necessary. – Transaction Log Resize actions based on the Transaction Log resource depend on support for vStorage in the underlying hypervisor technology. Because current versions of Hyper-V do not provide API support for vStorage, Intersight Workload Optimizer cannot support Transaction Log resize actions for database servers running on the Hyper-V platform.
Virtual Machine	■ **Resize** – Resize resource capacity 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

Entity:	Actions:
	VM. Before recommending this action, Intersight Workload Optimizer verifies that the VM's cluster can adequately support the new size. If the cluster is highly utilized, Intersight Workload Optimizer will recommend a move action, taking into consideration the capacity of the new cluster and compliance with existing placement policies. For hypervisor targets, Intersight Workload Optimizer can resize vCPU by changing the VM's socket or cores per socket count. For details, see "VCPU Scaling Controls" in the *User Guide*. – Resize resource reservation 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 – Resize resource limit 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. ■ **Move** Move a VM due to: – High resource utilization on VM or host – Excess IOPS or latency in VStorage – Workload placement violation – Underutilized host (move VM before suspending host) ■ **Move VM Storage (Volume)** Move a VM's volume due to excess utilization of the current datastore, or for more efficient utilization of datastores in the environment. ■ **Reconfigure** Change a VM's configuration to comply with a policy. For hypervisor targets, Intersight Workload Optimizer can reconfigure VMs that violate vCPU scaling policies. For details, see "VCPU Scaling Controls" in the *User Guide*. ■ **Reconfigure VM Storage** Reconfigure overutilized storage resources by adding VStorage capacity. For underutilized storage resources, remove VStorage capacity.

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the application server supply chain:

Entity Type	Commodity
Database Server	■ DBMem

Entity Type	Commodity
	The memory utilized by the database, as a of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Intersight Workload Optimizer can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM ■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ TransactionLog The utilization of the server's capacity for storage devoted to transaction logs Measured in Kilobytes (KB) ■ Cache Hit Rate The percentage of accesses that result in cache hits. Measured in a percentage of hits vs total attempts (%)
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM Measured in Kilobytes (KB) ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM Measured in IOPS ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)

Creating a Service User Account in Oracle

To collect data from the Oracle database, Intersight Workload Optimizer 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 Intersight Workload Optimizer needs:

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

Compute / Fabric Targets

A fabric target is a service that unites compute, network and storage access into a cohesive system. When you connect Intersight Workload Optimizer 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, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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

Intersight Workload Optimizer monitors the following resources for the fabric supply chain:

Entity Type	Commodity
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU)

Entity Type	Commodity
	The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Blade	■ Net The utilization of data through the Blade's network adapters Measured in Kilobytes per second (KB/s) ■ Treated as a Physical Machine of the underlying Hypervisor (see below) CPU, Mem, etc.
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB) ■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
I/O Module	■ NetThroughput Rate of message delivery over a port Measured in Megabits per second (Mb/s)
Switch	■ NetThroughput Rate of message delivery over a port Measured in Mb/s ■ PortChannel

Entity Type	Commodity
	Amalgamation of ports with a shared net throughput and utilization Measured in Mb/s

Cisco UCS Manager

Intersight Workload Optimizer supports UCS Manager 2.2+.

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. Intersight Workload Optimizer discovers these targets automatically.

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 Intersight Workload Optimizer enables automation at the hardware level, including automated provisioning of hosts.

Claiming UCS Targets

If your installation of Cisco Intersight has already claimed your UCS device, then Intersight Workload Optimizer discovers the UCS environment automatically.

To claim a new UCS device, select the **Compute / Fabric** category and choose the type of device you want for a target. Then provide the following:

- Device ID:

  Enter the applicable Device ID. Endpoint devices connect to the Cisco Intersight portal through a Device Connector that is embedded in the management controller (Management VM for Cisco UCS Director) of each system. The Device Connector provides a secure way for connected devices to send information and receive control instructions from the Cisco Intersight portal, using a secure Internet connection.
- Claim Code:

  The Claim Code authorizes your access. You can find this code in the Device Connector.
- Click **Claim**:

  After you provide the information, click **Claim**. You can see the status of your claimed target in the **Targets** tab.

The following table provides the format of the device ID and the device connector location:

Targets:	Device ID Format and Example:	Device Connector Location:
Standalone UCS Server	Serial Number Example: NGTR12345	From **Admin > Device Connector** in Cisco IMC
Intersight Managed Domain **Feature Preview**	Serial ID of the primary or subordinate FIs in this format: Serial number of FI-A or Serial number of FI-B	Redfish based device connector
Cisco UCS Manager **Feature Preview**	Serial ID of the primary and subordinate FIs in this format: Serial number of FI-A & Serial number of FI-B Example: [SAL1924GKV6&SAL1913CJ7V]	From **Admin > Device Connector** in Cisco IMC

Entity Mapping

Intersight Workload Optimizer Mapping	UCS
Host	Server / Blade / Rack Unit

Intersight Workload Optimizer Mapping	UCS
Chassis	Chassis
Datacenter	Datacenter
IO Module	IO Module
Switch	Fabric Interconnect
Network	Network

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

Intersight Workload Optimizer recommends actions for the various entities of the UCS Fabric Network as follows:

Entity Type	Action
Physical Machines	■ Start Physical Machine ■ Provision Physical Machine ■ Suspend Physical Machine
Chassis	■ Provision New Chassis
Fabric Interconnect	■ Add Port to Port Channel ■ Remove Port from Port Channel ■ Add Port
DPod (if Network Flow target is present)	■ Provision new DPod

Monitored Resources

Intersight Workload Optimizer monitors the following commodities of the UCS target:

Entity Type	Commodity
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space

Entity Type	Commodity
	Measured in Kilobytes (KB) ■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Chassis	■ Power Electricity being consumed by the Chassis Measured in Watts (W) ■ 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	■ NetThroughput Rate of message delivery over a port Measured in Megabits per second (Mb/s)
Switch	■ NetThroughput Rate of message delivery over a port Measured in Mb/s ■ PortChannel Amalgamation of ports with a shared net throughput and utilization Measured in Mb/s
DPod (if Network Flow target is present)	■ Memory (Mem) The utilization of the DPod's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the DPod's CPU reserved or in use Measured in Megahertz (MHz) ■ Storage The utilization of the storage attached to the DPod Measured in Kilobytes (KB) ■ Flow The utilization of the network flow capacity utilized by the DPod. This is divided into Flow1 (Low Cost) and Flow2 (Medium Cost) utilization Measured in Kilobytes per second (KB/s)

HPE OneView

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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 Intersight Workload Optimizer enables automation at the hardware level, including automated provisioning of hosts.

Prerequisites

- A service account Intersight Workload Optimizer 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.

  Intersight Workload Optimizer 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 Intersight Workload Optimizer will use to connect to the HPE OneView target.

  specify the IP address and credentials for HPE OneView. Intersight Workload Optimizer 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

Intersight Workload Optimizer recommends actions for the various entities of the HPE OneView Fabric Network as follows:

Entity Type	Action
Virtual Machines	■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Move Virtual Machine Storage ■ Reconfigure Storage ■ Reconfigure Virtual Machine ■ Suspend VM ■ Provision VM
Physical Machines	■ Start Physical Machine ■ Provision Physical Machine ■ Suspend Physical Machine
Fabric Interconnect	■ Add Port to Port Channel ■ Remove Port from Port Channel ■ Add Port
DPod (if Network Flow target is present)	■ Provision new DPod

Monitored Resources

Intersight Workload Optimizer monitors the following commodities of the HPE OneView target:

Entity Type	Commodity
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO

Entity Type	Commodity
	The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB) ■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
I/O Module	■ NetThroughput Rate of message delivery over a port Measured in Megabits per second (Mb/s)
Switch	■ NetThroughput Rate of message delivery over a port Measured in Mb/s ■ PortChannel Amalgamation of ports with a shared net throughput and utilization Measured in Mb/s
DPod (if Network Flow target is present)	■ Memory (Mem) The utilization of the DPod's memory reserved or in use

Entity Type	Commodity
	Measured in Kilobytes (KB) ■ CPU The utilization of the DPod's CPU reserved or in use Measured in Megahertz (MHz) ■ Storage The utilization of the storage attached to the DPod Measured in Kilobytes (KB) ■ Flow The utilization of the network flow capacity utilized by the DPod. This is divided into Flow1 (Low Cost) and Flow2 (Medium Cost) utilization Measured in Kilobytes per second (KB/s)

Application Performance Management (APM)

For APM, Intersight Workload Optimizer supports Cisco AppDynamics targets. These targets add Business Application, Business Transaction, Service, Application Component, and Database entities to the supply chain. To see how these entities map to the AppDynamics nomenclature, see Entity Mapping *(on page 83)*.

Cisco AppDynamics

NOTE:
This type of target can run as SaaS or in on-prem datacenters. When you claim the target, you can choose to turn ON or turn OFF **Connect through an Intersight Assist** as follows:

- If the target runs as SaaS:

  Turn OFF **Connect through an Intersight Assist**.

  You should be aware that for earlier versions of Intersight Workload Optimizer, to claim an AppDynamics target running as SaaS you were required to specify an Intersight Assist. If you claimed your target through an Assist, you can reclaim that target *without* using the Intersight Assist. To do that you must first delete the claimed target, and then claim the target anew with **Claim through an Intersight Assist** in the OFF position.
- If the target runs in an on-prem datacenter:

  Turn ON **Connect through an Intersight Assist**.

  To establish communication between this target on the datacenter and Intersight Workload Optimizer, you must:

  - Install an Intersight Assist appliance in the on-prem datacenter. The AppDynamics target must be accessible to the Intersight Assist appliance.
  - Connect the Intersight Assist instance with Cisco Intersight.
  - Log in to Cisco Intersight and claim the Intersight Assist instance as a target.
  - Claim the AppDynamics target with **Connect through an Intersight Assist** in the ON position.

    Intersight Assist provides a secure way for on-prem targets to send information to and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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

The Intersight Workload Optimizer integration with AppDynamics provides a full-stack view of your environment, from application to physical hardware. With information obtained from AppDynamics, Intersight Workload Optimizer 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 5000 AppDynamics nodes within the default collection period. Larger AppDynamics environments can take longer than one cycle to collect complete data.

Prerequisites

- A valid AppDynamics user account.

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

  NOTE:
  In newer versions of AppDynamics where these roles are available, they should be used instead:
  - Applications and Dashboards Viewer
  - DB Monitoring User
  - Server Monitoring

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

AppDynamics Database Servers

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

- Enable enhanced metric collection.

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

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

  For VMware hosts, you must install VMware Tools on the target VMs.
- Ensure that the database name in AppDynamics is resolvable to an IP address by the Intersight Workload Optimizer instance.

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

Entity Mapping

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

AppDynamics	Intersight Workload Optimizer
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

Claiming an AppDynamics Target

NOTE:
It is possible to monitor certain applications or database servers with both AppDynamics and Intersight Workload Optimizer, 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 Intersight Workload Optimizer application target.

To claim an AppDynamics instance as a target, specify:

- Connect through an Intersight Assist

  Whether to claim the target via an Intersight Assist instance.

  If your AppDynamics is deployed in your datacenter, then you must turn this ON and use an Intersight Assist to establish the connection with that target.

  If the target is a SaaS-based AppDynamics instance, then you should turn this option OFF.
- Intersight Assist

  The Intersight Assist instance that you will use to claim this AppDynamics target.

  To provide this setting you must turn on **Connect through an Intersight Assist**. You must also have already claimed at least one Intersight Assist instance.
- 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 @`
- Secure Connection

  When checked, Intersight Workload Optimizer 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.

For more information about creating API client users, see the AppDynamics Documentation.

Actions

NOTE:
The specific actions that Intersight Workload Optimizer recommends can differ, depending on the processes that Intersight Workload Optimizer discovers.

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

Intersight Workload Optimizer recommends actions for the AppDynamics supply chain as follows.

Entity Type	Action
Database Server	■ 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 Intersight Workload Optimizer 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 or Connections. The target does not discover DB Memory or Connections.
- MongoDB:

  For MongoDB database servers, analysis does not generate resize actions for DB Memory 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 Intersight Workload Optimizer does not display metrics for the VM.

Intersight Workload Optimizer monitors the following resources for the AppDynamics supply chain:

Entity Type	Commodity
Business Transaction	■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Transactions The utilization of the allocated transactions per second for the given business transaction Measured in transactions per second
Business Applications	■ Response Time

Entity Type	Commodity
	The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Transactions The utilization of the allocated transactions per second for the given virtual application Measured in transactions per second
Service	■ Response Time Response time (in milliseconds) for the given service For Kubernetes, this is the desired *weighted average* response time of all Application Component replicas associated with a Service ■ Transactions Transactions per second for the given service For Kubernetes, the maximum number of transactions per second that each Application Component replica can handle.
Application Component	■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) **NOTE:** This commodity is collected for Java, .NET, and Node.js applications only. ■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) **NOTE:** This commodity is collected for Java, .NET, and Node.js applications only. ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second ■ Heap The utilization of the application component's heap Measured in Kilobytes (KB) **NOTE:** This commodity is collected for Java, .NET, and Node.js applications only. ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ Remaining Garbage Collection Capacity The percentage of server uptime spent not performing garbage collection Measured in uptime (%) ■ Threads The utilization of the server's thread capacity Measured in number of Threads
Database Server	■ Virtual Memory (VMem)

Entity Type	Commodity
	The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) **NOTE:** Requires a machine agent present, and database hardware monitoring to be enabled. ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) **NOTE:** Requires a machine agent present, and database hardware monitoring to be enabled. ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second **NOTE:** For Microsoft SQL Server, MySQL, and Oracle databases only. ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections **NOTE:** For Mongo databases only. ■ Transaction Log The utilization of the server's capacity for storage devoted to transaction logs Measured in Kilobytes (KB) **NOTE:** For Microsoft SQL Server databases only. ■ DB Cache Hit Rate The percentage of accesses that result in cache hits. Measured as a percentage of hits vs total attempts (%) **NOTE:** For Microsoft SQL Server and Oracle databases only.
Virtual Machine	■ Virtual CPU (vCPU) The utilization of the VCPU of the virtual machine. Measured in %. ■ Virtual Memory (vMem) The utilization of the VMEM of the virtual machine. Measured in Kilobytes (KB).

New Relic

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

Prerequisites

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

Entity Mapping

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

New Relic Term	Intersight Workload Optimizer 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

Intersight Workload Optimizer 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

Intersight Workload Optimizer supports the following Database types and commodities:

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

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

Claiming a New Relic Target

NOTE:
If an application is monitored by New Relic, do not add it as a separate Intersight Workload Optimizer application target.

To claim New Relic as a target, specify:

- Name

  The target name that displays in the user interface.
- REST API Key

  The REST API Key *provided by the New Relic platform*. For more information, see New Relic API Keys.
- Account ID

  The New Relic Account ID.
- 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.
- Region

  If checked, Intersight Workload Optimizer will use the EU API endpoints.

Actions

NOTE:
The specific actions that Intersight Workload Optimizer recommends can differ, depending on the processes that Intersight Workload Optimizer discovers.

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

Intersight Workload Optimizer recommends actions for the New Relic supply chain as follows.

Entity Type	Action
Application Component	■ 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 Intersight Workload Optimizer does not display metrics for the VM.

Intersight Workload Optimizer monitors the following resources for the New Relic supply chain:

Entity Type	Commodity
Application Component	■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM Measured in Megahertz (MHz) ■ Virtual Memory (VMem)

Entity Type	Commodity
	The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Transactions Intersight Workload Optimizer discovers *key* transactions – transactions that the user marked as *key* to the application The utilization of the allocated transactions per second for the given entity Measured in transactions per second ■ Heap The utilization of the application server's heap Measured in Kilobytes (KB) ■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ Remaining Garbage Collection Capacity The percentage of uptime that is *not* spent on garbage collection. Measured in uptime (%) ■ Threads The utilization of the server's thread capacity Measured in number of Threads
Database	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM Measured in Kilobytes (KB) ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second. ■ DBMem The memory utilized by the database, as a percentage of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Intersight Workload Optimizer can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM. ■ Connections The utilization of the connection capacity. Only applicable to database servers Measured in Connections ■ DB Cache Hit Rate The percentage of accesses that result in cache hits. Measured as a percentage of hits vs total attempts (%)
Business Transaction	■ Response Time The utilization of the server's allocated response time. Measured in milliseconds (ms). ■ Transactions

Entity Type	Commodity
	The utilization of the allocated transactions per second for the given entity Measured in transactions per second.
Service	■ Response Time The utilization of the server's allocated response time. Measured in milliseconds (ms). ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second.
Virtual Machine	■ Virtual CPU (vCPU) The utilization of the VCPU of the virtual machine. Measured in %. ■ Virtual Memory (vMEM) The utilization of the VMEM of the virtual machine. Measured in Kilobytes (KB).

Dynatrace

NOTE:
This type of target can run as SaaS or in on-prem datacenters. When you claim the target, you can choose to turn ON or turn OFF **Connect through an Intersight Assist** as follows:

- If the target runs as SaaS:

  Turn OFF **Connect through an Intersight Assist**.
- If the target runs in an on-prem datacenter:

  Turn ON **Connect through an Intersight Assist**.

  To establish communication between this target on the datacenter and Intersight Workload Optimizer, you must:
  - Install an Intersight Assist appliance in the on-prem datacenter. The target you are configuring must be accessible to the Intersight Assist appliance.
  - Connect the Intersight Assist instance with Cisco Intersight.
  - Log in to Cisco Intersight and claim the Intersight Assist instance as a target.
  - Claim the target you are configuring with **Connect through an Intersight Assist** in the ON position.

    Intersight Assist provides a secure way for on-prem targets to send information to and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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

Prerequisites

- A Dynatrace Server instance

  This instance must be configured to monitor applications running in your environment.

  Intersight Workload Optimizer supports both SaaS and on-prem Dynatrace server installations.
- Managed VMs that host applications managed by Dynatrace

  For Intersight Workload Optimizer to discover applications through Dynatrace, the applications must be running on VMs in your environment. Also, these VMs must be managed by Intersight Workload Optimizer targets such as hypervisors or public cloud targets.
- An API access token with the proper scopes

Intersight Workload Optimizer uses an API token to authenticate its calls to the Dynatrace API. Use a generic token with these scopes:

Intersight Workload Optimizer Functionality	Required Permissions
Monitoring	– API V1 scopes: • `Access problem and event feed, metrics, and topology` – API V2 scopes: • `Read entities` • `Read metrics`

NOTE:
If you are updating to Intersight Workload Optimizer version or later, from a version that is earlier than , 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.

- The Environment ID

  To claim a Dynatrace target, you must know the Environment ID for the given Dynatrace installation. According to the Dynatrace documentation, you can identify the Environment ID in these ways:

  – SaaS-based Dynatrace Server:

    The Environment ID is the first part of the Dynatrace environment URL. For example, for the environment `https://abc123a.live.dynatrace.com`, the Environment ID is `abc123a`

  – On-prem Dynatrace Server:

    The Environment ID is the string after `/e/` in the Dynatrace environment URL. For example, for the environment address `https://managed-cluster/e/abc123a`, the Environment ID is `abc123a`

  For more information, see the Dynatrace documentation at `https://www.dynatrace.com/support/help/get-started/monitoring-environment/environment-id/`

Entity Mapping

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

Dynatrace Naming	Intersight Workload Optimizer Entity
Application	Business Application **NOTE:** For Dynatrace Applications, Intersight Workload Optimizer 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

Claiming a Dynatrace Target

NOTE:
You can manage certain applications or database servers with both Dynatrace and Intersight Workload Optimizer. You should avoid such a configuration because it can cause Intersight Workload Optimizer to generate duplicate entities in the market.

If you manage an application via a Dynatrace server, and you configure that Dynatrace server as a Intersight Workload Optimizer target, then be sure you have *not* added that application as a separate application target in Intersight Workload Optimizer.

On-Prem Dynatrace Target:

To claim an *on-prem* Dynatrace server instance as a target, specify:

- Connect through an Intersight Assist

  Whether to claim the target via an Intersight Assist instance.

  If your Dynatrace server is deployed in your datacenter, then you must turn this ON and use an Intersight Assist to establish the connection with that target.
- Intersight Assist

  The Intersight Assist instance that you will use to claim this Dynatrace target.

  To provide this setting you must turn on **Connect through an Intersight Assist**. You must also have already claimed at least one Intersight Assist instance.
- Hostname or IP Address

  For an on-prem installation of DynaTrace, give the host name or IP for your Dynatrace server. For example, `10.10.12.34`.
- Environment ID

  The unique string that identifies the environment this Dynatrace target manages.
- API Key

  The token that Intersight Workload Optimizer 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.

SaaS-Based Dynatrace Target:

To claim a *SaaS-based* Dynatrace server instance as a target, specify:

- Connect through an Intersight Assist

  Whether to claim the target via an Intersight Assist instance.

  If the target is a SaaS-based Dynatrace server, turn this option OFF.
- Environment ID

  The unique string that identifies the environment this Dynatrace target manages.
- API Key

  The token that Intersight Workload Optimizer 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.

Actions

Entity:	Actions:
Application Component	■ **Resize** – Heap Recommended, only.

Monitored Resources

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

For Database Server applications, Intersight Workload Optimizer 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 Intersight Workload Optimizer does not display metrics for the VM.

Intersight Workload Optimizer monitors the following resources for the Dynatrace supply chain:

Entity Type	Commodity
Business Application	■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second
Service	■ Response Time The utilization of the server's allocated response time Measured in Milliseconds (ms) ■ Transactions The utilization of the allocated transactions per second for the given entity Measured in transactions per second
Application Component	■ Virtual CPU (vCPU) The utilization of the VCPU for the given entity. Measured in %. ■ Virtual Memory (vMem) The utilization of the VMEM for the given entity. Measured in Kilobytes (KB). ■ Remaining Garbage Collection Capacity The percentage of server uptime spent not performing garbage collection Measured in uptime (%) **NOTE:** This commodity is for Java applications only. ■ Heap The utilization of the application server's heap Measured in Kilobytes (KB) **NOTE:** This commodity is collected for Java applications only.
Database Server	■ Virtual CPU (vCPU) The utilization of the VCPU for the given entity. Measured in %.

Entity Type	Commodity
	■ Virtual Memory (vMem) The utilization of the VMEM for the given entity. Measured in Kilobytes (KB). ■ DBMem For Microsoft SQL and MySQL only. The memory utilized by the database, as a percentage of the memory capacity that is allocated to the database. Note that this resource is more accurate than the VMem resource on the hosting VM. With this resource, Intersight Workload Optimizer can drive resize and move actions based on the memory consumed by the database, not the memory consumed by the VM. ■ DB Cache Hit Rate For Microsoft SQL only. The percentage of accesses that result in cache hits. Measured as a percentage of hits vs total attempts (%) ■ Transactions For Microsoft SQL only. The utilization of the allocated transactions per second for the given entity Measured in transactions per second
Container	■ Virtual CPU (vCPU) The utilization of the VCPU for the given entity. Measured in %. ■ Virtual Memory (vMem) The utilization of the VMEM for the given entity. Measured in Kilobytes (KB).
Virtual Machine	■ Virtual CPU (vCPU) The utilization of the VCPU for the given entity. Measured in %. ■ Virtual Memory (vMem) The utilization of the VMEM for the given entity. Measured in Kilobytes (KB).

Hyperconverged Targets

A hyperconverged target is a service that unites compute, network and storage access into a cohesive system. When you connect Intersight Workload Optimizer 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, Intersight Workload Optimizer 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

Intersight Workload Optimizer 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	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Blade	■ Net The utilization of data through the Blade's network adapters Measured in Kilobytes per second (KB/s) ■ Treated as a Physical Machine of the underlying Hypervisor (see below)

Entity Type	Commodity
	CPU, Mem, etc.
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB) ■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
I/O Module	■ NetThroughput Rate of message delivery over a port Measured in Megabits per second (Mb/s)
Switch	■ NetThroughput Rate of message delivery over a port Measured in Mb/s ■ PortChannel Amalgamation of ports with a shared net throughput and utilization Measured in Mb/s
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second

Entity Type	Commodity
	NOTE: When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)

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. Intersight Workload Optimizer discovers these targets automatically.

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

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. Intersight Workload Optimizer will not recommend actions for these VMs.

Claiming HyperFlex Targets

If your installation of Cisco Intersight has already claimed your HyperFlex device, then Intersight Workload Optimizer discovers the HyperFlex environment automatically.

To claim a new HyperFlex device, select the **Hyperconverged** category and choose the type of device you want for a target. Then provide the following:

- Device ID:

Enter the applicable Device ID. Endpoint devices connect to the Cisco Intersight portal through a Device Connector that is embedded in the management controller (Management VM for Cisco UCS Director) of each system. The Device Connector provides a secure way for connected devices to send information and receive control instructions from the Cisco Intersight portal, using a secure Internet connection.

  - The device ID format is the Cluster UUID. For example, xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
  - Locate the device connector location via **HyperFlex Connect UI > Settings > Device Connector** in Cisco HyperFlex.

- Claim Code:

  The Claim Code authorizes your access. You can find this code in the Device Connector.

- Click **Claim**:

  After you provide the information, click **Claim**. You can see the status of your claimed target in the **Targets** tab.

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, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in HyperFlex to those used in Intersight Workload Optimizer:

HyperFlex Name	Intersight Workload Optimizer Entity
Volume	Storage
HX Cluster	Disk Array

Supported Actions

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

NOTE:
For this target, Intersight Workload Optimizer 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

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS)

Entity Type	Commodity
	The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)

Nutanix Acropolis

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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.

Intersight Workload Optimizer supports management of Nutanix fabrics, where the supply chain treats a Nutanix Storage Pool as a disk array. Intersight Workload Optimizer recognizes Nutanix storage tiers when calculating placement of VMs and VStorage. In addition, Intersight Workload Optimizer 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 Acropolis 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, Intersight Workload Optimizer 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 104)*.

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 Acropolis 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

Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 104)*.

Adding Nutanix Targets

NOTE:
This describes how to add a Nutanix cluster to Intersight Workload Optimizer 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 105)*.

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

  WhenIntersight Workload Optimizer will use a secure connection.
- Username/Password

  Credentials for an account on the Nutanix cluster with sufficient privileges.

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

Nutanix Name	Intersight Workload Optimizer Entity
Container	Storage
Storage Pool	Disk Array
Nutanix Cluster	Storage Controller

Supported Actions

For each discovered entity, Intersight Workload Optimizer 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:** Intersight Workload Optimizer can automate VMotion to hosts, but for storage moves on Nutanix Intersight Workload Optimizer only supports the Recommend action mode.	
Datastore ("Storage")	Provision, Resize Up, Resize Down, Suspend	Move
Disk Array		
Storage Controller		Provision

Monitored Resources

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount

Entity Type	Commodity
	The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	**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. ■ CPU The utilization of the Storage Controller's allocated CPU Measured in Megahertz (MHz) ■ Storage Amount The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller. Measured in Megabytes (MB)

Pinning Nutanix Controller VMs in Generic Hypervisor Mode

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

For a cluster in Generic Hypervisor mode (using vCenter or Hyper-V hypervisors), you must use Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer navigate to **IWO > More > 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 Intersight Workload Optimizer, navigate to **IWO > More > Settings > Policy**, and create a scoped automation policy for VMs.
   - Set the scope to the group you made
   - Disable moves for this group
   - 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 Intersight Workload Optimizer to hypervisor targets in your environment, Intersight Workload Optimizer assures application performance by utilizing these resources as efficiently as possible.

Once connected to a hypervisor target, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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.

Intersight Workload Optimizer 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, Intersight Workload Optimizer 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 Intersight Workload Optimizer extends the supply chain analysis into the components that make up the disk array. For more information, see Storage Manager Targets *(on page 121)*.

Actions

Intersight Workload Optimizer 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	■ Provision additional resources (VMem, VCPU) ■ Move Virtual Machine ■ Move Virtual Machine Storage

Entity Type	Action
	■ Reconfigure Storage ■ Reconfigure Virtual Machine ■ Suspend VM ■ Provision VM
Physical Machines	■ Start Physical Machine ■ Provision Physical Machine ■ Suspend Physical Machine
Storage	■ Start Storage ■ Provision Storage ■ Suspend Storage ■ Move (only with Storage Targets configured) ■ Resize (only with Storage Targets configured)
Consumer Virtual Datacenters	■ Resize Consumer vDC ■ Provision Consumer vDC

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space

Entity Type	Commodity
	Measured in Kilobytes (KB) ■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Datacenter	**NOTE:** For datacenter entities, Intersight Workload Optimizer does not monitor resources directly from the datacenter, but from the physical machines in the datacenter. ■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB) ■ Balloon

Entity Type	Commodity
	The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Kilobytes (KB)
Provider Virtual Datacenter	■ Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz) ■ Storage The utilization of the storage attached to the Provider vDC. Measured in Kilobytes (KB)
Consumer Virtual Datacenter	■ Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz) ■ Storage The utilization of the storage attached to the Consumer vDC. Measured in Kilobytes (KB)

Microsoft Hyper-V

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

If you have a small number of Hyper-V hosts in your environment, you can add them individually as Intersight Workload Optimizer 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 Intersight Workload Optimizer automatically add the other members of that cluster.

For accurate SMB storage calculations, Intersight Workload Optimizer requires a VMM target.

Prerequisites

- Create a user account that Intersight Workload Optimizer can use to connect to your Hyper-V servers. See Creating a Service User Account in Hyper-V *(on page 111)*
- Configure remote management on each Hyper-V server. Refer to Enabling Windows Remote Management *(on page 153)*
- 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, Intersight Workload Optimizer 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

  Intersight Workload Optimizer 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 154)*).
- Username

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

  Password for account used.

Exporting Hyper-V Virtual Machines

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

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

Intersight Workload Optimizer 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, Intersight Workload Optimizer can execute or recommend certain actions, as outlined below.

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

Entity Type	Can Be Automated	Recommendations Only
Physical Machine	Start, Suspend	Terminate, Provision
Storage		Provision

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB)
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second

Entity Type	Commodity
	NOTE: When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Datacenter	**NOTE:** For datacenter entities, Intersight Workload Optimizer does not monitor resources directly from the datacenter, but from the physical machines in the datacenter. ■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB)

Creating A Service User Account

The service account Intersight Workload Optimizer 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 Intersight Workload Optimizer 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:
Intersight Workload Optimizer 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

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

VMware vCenter Server provides a centralized management platform for VMware hypervisors. To manage your VMware environment with Intersight Workload Optimizer, you specify a vCenter Server instance as a target. Intersight Workload Optimizer 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 Intersight Workload Optimizer uses to connect to your vCenter and execute actions. For information about setting permissions for this account, see Creating a User Account In vCenter *(on page 116)*.

General Considerations

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

- **Linked vCenters**:

  For linked vCenters, you must add each vCenter separately so Intersight Workload Optimizer can communicate with each vCenter through a separate API endpoint.

- **Restricting Intersight Workload Optimizer Access to Specific Clusters**:

  When you add a vCenter target, Intersight Workload Optimizer discovers all of the connected entities that are visible, based on the target account that it uses to connect to the vCenter 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 Intersight Workload Optimizer account to `No access` for the given entities.

- **Shared Datastores**:

  If you add more than one vCenter 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 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 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 Intersight Workload Optimizer is assigned has the `Profile-driven storage view` permission enabled. This permission is *disabled* in the built-in `readonly` role.

Claiming vCenter Targets

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

- Address

  The name or IP address of the vCenter server.

- Username/Password

  Credentials for the user account Intersight Workload Optimizer can use to connect to the vCenter Server. Include the domain if required (*<domain>**<username>*).

- Enable Datastore Browsing

  Enabling datastore browsing enables Intersight Workload Optimizer to discover wasted storage.

- Enable Guest Metrics

  Collecting advanced guest memory metrics can increase the accuracy of the VMEM data that Intersight Workload Optimizer uses for analysis of virtual machines. To enable guest metrics, ensure the following:

  - VMware Tools is installed and running on the target VMs
  - The **Hypervisor VMEM for Resize** vCenter option is active for the VMs discovered by the target.
  - The user account has the `Performance.Modify Intervals` performance privilege

  For more information, see "Hypervisor VMEM for Resize" in the *User Guide* and vCenter Performance Privileges.

vCenter Imported Settings

In addition to discovering entities managed by the hypervisor, Intersight Workload Optimizer 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 117)*).

NOTE: Intersight Workload Optimizer 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 Intersight Workload Optimizer to include cross vCenter vMotions in its recommendations.

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

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

   To merge datacenters choose Datacenter, to merge Host clusters choose Cluster, and for storage choose StorageCluster.
4. Choose the specific datacenters or clusters to merge in this policy, then click **Select**.
5. Click **Save Policy**.

NOTE:
Since Intersight Workload Optimizer 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 Intersight Workload Optimizer will not initiate VSS → VDS vMotions, vSphere may do so. If this happens, Intersight Workload Optimizer 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, Suspend, Storage Move, Resize Down, Resize Up	Terminate, Provision, Reconfigure
Physical Machine	Start, Suspend	Terminate, Provision
Storage		Provision

Monitored Resources

Intersight Workload Optimizer monitors the following resources for the hypervisor supply chain:

Entity Type	Commodity
Virtual Machine	■ Virtual Memory (VMem) The utilization of the VMem allocated to the hosting VM ■ Virtual CPU (VCPU) The utilization of the VCPU allocated to the hosting VM ■ Virtual Storage (VStorage) The utilization of the virtual storage capacity allocated for the VM ■ Storage Access Operations Per Second (IOPS) The utilization of IOPS allocated for the VStorage on the VM ■ Latency The utilization of latency allocated for the VStorage on the VM Measured in milliseconds (ms)
Host	■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net

Entity Type	Commodity
	The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s) ■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB) ■ Balloon The utilization of shared memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Megahertz (MHz)
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Datacenter	**NOTE:** For datacenter entities, Intersight Workload Optimizer does not monitor resources directly from the datacenter, but from the physical machines in the datacenter. ■ Memory (Mem) The utilization of the PM's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the PM's CPU reserved or in use Measured in Megahertz (MHz) ■ IO The utilization of the PM's IO adapters Measured in Kilobytes per second (KB/s) ■ Net The utilization of data through the PM's network adapters Measured in Kilobytes per second (KB/s)

Entity Type	Commodity
	■ Swap The utilization of the PM's swap space Measured in Kilobytes (KB) ■ Balloon The utilization of shared of memory among VMs running on the host. ESX-only Measured in Kilobytes (KB) ■ CPU Ready The utilization of the PM's allocated ready queue capacity (measured in Kbytes) that is in use, for 1, 2, and 4 CPU ready queues. ESX-only Measured in Kilobytes (KB)
Provider Virtual Datacenter	■ Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz) ■ Storage The utilization of the storage attached to the Provider vDC. Measured in Kilobytes (KB)
Consumer Virtual Datacenter	■ Memory (Mem) The utilization of the Datacenter's memory reserved or in use Measured in Kilobytes (KB) ■ CPU The utilization of the Datacenter's CPU reserved or in use Measured in Megahertz (MHz) ■ Storage The utilization of the storage attached to the Consumer vDC. Measured in Kilobytes (KB)

Creating A Service User Account In vCenter

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

vCenter Permissions

Intersight Workload Optimizer Functionality	Required Permissions
Monitoring	■ 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. ■ Requirement to monitor VSAN and storage profiles

Intersight Workload Optimizer Functionality	Required Permissions
	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 *disabled* in the built-in `readonly` role.
Recommend Actions	■ 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	■ `Datastore > Browse Datastore`
Execute VM Move	■ `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	■ `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	■ `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	■ `Global > Global tag` You must also open ports `10443` and `7443` on the target server

Other Information Imported from vCenter

In addition to discovering entities managed by the vSphere hypervisors and their resources, Intersight Workload Optimizer:

- 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 Intersight Workload Optimizer.

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

Intersight Workload Optimizer supports the ServiceNow orchestrator target.

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

ServiceNow

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

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 Intersight Workload Optimizer 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 Intersight Workload Optimizer will use to connect to the ServiceNow instance
- Password

  Password for the account Intersight Workload Optimizer will use to connect to the ServiceNow instance
- Client ID

  The Client ID Intersight Workload Optimizer will use if `Use oAuth` is checked
- Client Secret

  The password Intersight Workload Optimizer will use if `Use oAuth` is checked
- Port

Port used to access the ServiceNow Instance

- Use oAuth

  When checked, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer Actions for ServiceNow documentation.

Storage Targets

Adding a storage Target enables Intersight Workload Optimizer to connect to your storage subsystem through a native or SMI-S provider API. Intersight Workload Optimizer 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, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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, Intersight Workload Optimizerwill 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 Intersight Workload Optimizer 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

Intersight Workload Optimizer 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

Intersight Workload Optimizer recommends actions for storage targets as follows.

NOTE:
This is a general list of actions for storage managed by storage controllers. The actions that Intersight Workload Optimizer can recommend or automate depend on the actual technology. For example, Intersight Workload Optimizer 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	■ Start Storage ■ Provision Storage ■ Suspend Storage ■ Move (only with Storage Targets configured) ■ Resize (only with Storage Targets configured)
Disk Arrays	■ Provision Disk Array ■ Start Disk Array ■ Suspend Disk Array ■ Move Disk Array (for NetApp Cluster-Mode, only) ■ Move Virtual Machine ■ Move Datastore
Storage Controller	■ Provision Storage Controller (recommendation only)

Monitored Resources

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB)

Entity Type	Commodity
	■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	**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. ■ CPU The utilization of the Storage Controller's allocated CPU Measured in Megahertz (MHz) ■ Storage Amount The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller. Measured in Megabytes (MB)

Dell EMC SC Series

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

Intersight Workload Optimizer supports the management of Dell SC Series (Compellent) disk arrays and storage controllers. Intersight Workload Optimizer connects through the Dell Enterprise Manager and performs management as a client of the Enterprise Manager Data Collector.

The Dell Enterprise Manager is a management service that provides administration, management, and monitoring of multiple Storage Centers – Typically installed on a Windows VM.

When you specify a Dell Compellent target, you provide the IP address of the Dell Enterprise Manager. Intersight Workload Optimizer discovers the Compellent infrastructure through the SMI-S component which is typically installed as part of the Enterprise Manager.

NOTE: Before adding the Dell Compellent target to Intersight Workload Optimizer, confirm that the Storage Centers you want to manage show up in Dell Enterprise Manager (see "Storage Center Administration" in the *Dell Compellent Enterprise Manager Administrator's Guide*). The SMI-S user account must be able to access all of the Storage Centers. If you add or remove Storage Centers later, Intersight Workload Optimizer will detect the changes during its next discovery cycle.

Prerequisites

- Dell Enterprise Manager Data Collector Service 6.2 or higher
- Dell Compellent SMI-S Provider
- Storage Centers added to Dell Enterprise Manager

Setting Up the Dell Compellent SMI-S Provider

Your Dell Compellent storage environment must include an enabled Dell Compellent SMI-S Provider. Configure the SMI-S Provider as described in the "SMI-S" section of the *Dell Storage Manager Administrator's Guide*. The guide provides detailed steps to:

- Open the required ports on the server hosting the Enterprise Manager Data Collector.
- Enable SMI-S for the Data Collector.
- Add a user for SMI-S.
- If using HTTPS, associate the SSL certificate with the SMI-S Provider.

Adding Dell Compellent Targets

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

- Address

  The name or IP address of the Dell Enterprise Manager.

  By default, Enterprise Manager provides SMI-S data over port 5988 (HTTP). If your installation uses a different HTTP port for SMI-S, include the port number in the Address field. For HTTPS, do not include the port number. Instead, select the `Use Secure Connection` check box.
- Username/Password

  Credentials for the SMI-S user you added when setting up the SMI-S provider.
- Use Secure connection

  Select this option to connect to the target using a secure connection (HTTPS). Do not enter a port in the Address field if this option is selected.

After validating the new target, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in the Dell Enterprise Manager to those used in Intersight Workload Optimizer:

Dell Name	Intersight Workload Optimizer Entity
Storage Center	Storage Controller
Storage Type	Disk Array
Volume	Storage

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 121)*.

Supported Actions

Intersight Workload Optimizer supports the following actions for Dell Compellent entities:

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

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

Monitored Resources

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	**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. ■ CPU The utilization of the Storage Controller's allocated CPU

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

EMC VMAX

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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

NOTE:
Intersight Workload Optimizer does not utilize Unisphere. Data is collected exclusively from the SMI-S provider.

Intersight Workload Optimizer 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 Intersight Workload Optimizer can use to connect to the EMC SMI-S Provider (typically the default `admin` account)

Claiming VMAX Targets

To claim 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, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in EMC VMAX to those used in Intersight Workload Optimizer:

EMC VMAX Name	Intersight Workload Optimizer 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, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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.

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Logical Pool	■ Storage Amount The utilization of the logical pool's capacity. Measured in Megabytes (MB) ■ Storage Provisioned

Entity Type	Commodity
	The utilization of the logical pool's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the logical pool. Measured in Operations per second ■ Latency The utilization of latency on the logical pool. Measured in milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	■ Storage Amount The utilization of the storage controller's capacity. Measured in Megabytes (MB)

EMC XtremIO

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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).

Intersight Workload Optimizer 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 Intersight Workload Optimizer, 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 — Intersight Workload Optimizer cannot make recommendations to move datastores from one X-Brick to another. Additionally, the X-Brick has a fixed form factor — Intersight Workload Optimizer does not recommend resize actions for disk array or storage controller resources.

Intersight Workload Optimizer 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

  Intersight Workload Optimizer uses this account to connect to the XMS and execute commands through the XtremIO API.

Claiming 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, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in XtremIO to those used in Intersight Workload Optimizer:

XTremIO Name	Intersight Workload Optimizer 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 121)*.

Supported Actions

For each discovered entity, Intersight Workload Optimizer 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, Intersight Workload Optimizer excludes disks devoted to the VNX operating system.

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount

Entity Type	Commodity
	The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	**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. ■ CPU The utilization of the Storage Controller's allocated CPU Measured in Megahertz (MHz) ■ Storage Amount The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller. Measured in Megabytes (MB)

EMC ScaleIO

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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.

Intersight Workload Optimizer communicates with the EMC ScaleIO system via the REST API Gateway.

Prerequisites

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

Claiming EMC ScaleIO Targets

To claim 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, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in EMC ScaleIO to those used in Intersight Workload Optimizer:

EMC ScaleIO Name	Intersight Workload Optimizer Entity
Volume	Storage
Storage Pool	Disk Array
Protection Domain	Storage Controller

Supported Actions

For each discovered entity, Intersight Workload Optimizer 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

Entity Type	Can Be Automated	Recommendations only
Protection Domain		Provision (Cloning)

Monitored Resources

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	**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. ■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	■ Storage Amount The utilization of the storage controller's capacity. Measured in Megabytes (MB)

EMC VPLEX

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

Intersight Workload Optimizer supports management of EMC VPLEX virtual storage systems in a local configuration, via the VPLEX API. Currently, Intersight Workload Optimizer 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 Intersight Workload Optimizer.

Prerequisites

- VPLEX Management Server
- Hypervisor target supported by Intersight Workload Optimizer
- Storage target supported by Intersight Workload Optimizer

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

VPLEX Permissions

Intersight Workload Optimizer Functionality	Required Permissions
Monitoring	■ Service Account
Action Execution	■ Admin account

Claiming EMC VPLEX Targets

To claim 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. Intersight Workload Optimizer will use the enhanced visibility provided by VPLEX to make more intelligent storage decisions- for example, recommending storage vMotion between pools.

HPE 3PAR

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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

  NOTE:
  For discovery and monitoring, the Intersight Workload Optimizer service account must have the `Browse` permission on all monitored domains. To exclude domains from monitoring, the Intersight Workload Optimizer service account must have no permissions on those domains. For action execution, Intersight Workload Optimizer 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, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in HPE 3PAR to those used in Intersight Workload Optimizer:

HPE 3PAR Name	Intersight Workload Optimizer 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 121)*.

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, Intersight Workload Optimizer:

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

  In the Intersight Workload Optimizer user interface, these disk arrays do not host storage – They appear empty. Intersight Workload Optimizer 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. Intersight Workload Optimizer considers this single logical pool when it performs analysis – It can recommend moving storage into or out of the AO group. Also,

Intersight Workload Optimizer 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 Intersight Workload Optimizer user interface. The display name for this logical pool is the name of the AO Configuration.

Supported Actions

For each discovered entity, Intersight Workload Optimizer 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

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency

Entity Type	Commodity
	The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Logical Pool	■ Storage Amount The utilization of the logical pool's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the logical pool's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the logical pool. Measured in Operations per second ■ Latency The utilization of latency on the logical pool. Measured in milliseconds (ms)
Storage Controller	**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. ■ CPU The utilization of the Storage Controller's allocated CPU Measured in Megahertz (MHz) ■ Storage Amount The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller. Measured in Megabytes (MB)

NetApp

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

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. Intersight Workload Optimizer connects to these storage controllers to support NetApp targets in Cluster-Mode (C-Mode).

Prerequisites

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

Enabling TLS

Starting with version 5.4, by default Intersight Workload Optimizer 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 Intersight Workload Optimizer 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, Intersight Workload Optimizer 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 140)*

Claiming NetApp Targets

To claim 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 Intersight Workload Optimizer to use.

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

NetApp Name	Intersight Workload Optimizer 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 121)*.

Supported Actions

For each discovered entity, Intersight Workload Optimizer 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 Intersight Workload Optimizer 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

Intersight Workload Optimizer 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	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)

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

Restricted Service Accounts In NetApp

While Intersight Workload Optimizer 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 Intersight Workload OptimizerRole
   ```
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 StorageFlashArray

NOTE:
This target runs in on-prem datacenters. To establish communication between targets on the datacenter and Intersight Workload Optimizer, you must:

- Install an Intersight Assist appliance in the on-prem datacenter. The target service must be accessible to the Intersight Assist appliance.
- Connect the Intersight Assist instance with Cisco Intersight.
- Log in to Cisco Intersight and claim the Intersight Assist instance as a target.

Intersight Assist provides a secure way for connected targets to send information and receive control instructions from Intersight Workload Optimizer, using a secure internet connection. For more information, see the Cisco Intersight Assist Getting Started Guide.

Intersight Workload Optimizer 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, Intersight Workload Optimizer intelligently moves more demanding workloads to Flash-based datastores. Intersight Workload Optimizer analysis is also able to incorporate Pure Storage de-duplication and compression when recommending actions.

Prerequisites

- A service account Intersight Workload Optimizer 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.

Claiming Pure Storage Targets

To claim 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 Intersight Workload Optimizer 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, Intersight Workload Optimizer discovers the connected storage entities. This table compares terms used in Pure to those used in Intersight Workload Optimizer:

Pure Name	Intersight Workload Optimizer 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 121)*.

Supported Actions

For each discovered entity, Intersight Workload Optimizer 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

Intersight Workload Optimizer monitors the following storage resources:

Entity Type	Commodity
Storage	■ Storage Amount The utilization of the datastore's capacity Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the datastore's capacity, including overprovisioning.

Entity Type	Commodity
	Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the datastore Measured in Operations per second **NOTE:** When it generates actions, Intersight Workload Optimizer does not consider IOPS throttling that it discovers on storage entities. Analysis uses the IOPS it discovers on Logical Pool or Disk Array entities. ■ Latency The utilization of latency on the datastore Measured in Milliseconds (ms)
Disk Array	■ Storage Amount The utilization of the Disk Array's capacity. Measured in Megabytes (MB) ■ Storage Provisioned The utilization of the Disk Array's capacity, including overprovisioning. Measured in Megabytes (MB) ■ Storage Access Operations Per Second (IOPS) The summation of the read and write access operations per second on the disk array Measured in Operations per second ■ Latency The utilization of latency, computed from the latency of each device in the disk array. Measured in milliseconds (ms)
Storage Controller	**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. ■ CPU The utilization of the Storage Controller's allocated CPU Measured in Megahertz (MHz) ■ Storage Amount The utilization of the storage controller's capacity. The storage allocated to a storage controller is the total of all the physical space available to aggregates managed by that storage controller. Measured in Megabytes (MB)

Appendix — Target Configuration

This appendix contains topics that are related to configuring Intersight Workload Optimizer targets.

Cisco Unified Computing System

UCS Blade Provisioning

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

  Intersight Workload Optimizer 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, Intersight Workload Optimizer sets the Host Provision action to Recommend.

  For any hosts other than blade servers managed by UCS, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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, Intersight Workload Optimizer 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.

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

Enabling Collection of Memory Statistics: AWS

So Intersight Workload Optimizer can collect memory statistics in your AWS environment, you must set up your VMs to publish statistics via CloudWatch. 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 Intersight Workload Optimizer

  For Intersight Workload Optimizer 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 claiming an AWS target.

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 Intersight Workload Optimizer analysis and they do not show up in the user interface. Intersight Workload Optimizer 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.

Enabling Collection of Memory Statistics: Azure

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

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

- For Windows VMs

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

- For Linux VMs

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

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

GCP Target Service Account

This appendix guides you through the steps to create a valid service account that Intersight Workload Optimizer can use to connect with a GCP operational target. The access you grant to this service account determine the access Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer.

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 `IWOSaAccessProject.yaml`. Edit the file to have the following content:

```
title: "IWO Role: Min Accecss - Project"
description: "Minimal Required Permissions for IWO 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 `IWOSaAccessProject.yaml` with the path to your yaml file:

```
$ gcloud iam roles create $CUSTOM_ROLE_NAME \
    --project=$PROJECT_ID  \
    --file=IWOSaAccessProject.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 `IWOSaAccessOrg.yaml`. Edit the file to have the following content:

```
title: "IWO Role: Access - Organization"
description: "Minimal Required Permissions for IWO 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 `IWOSaAccessOrg.yaml` with the path to your yaml file:

```
$ gcloud iam roles create $CUSTOM_ROLE_NAME \
    --organization=$ORGANIZATION_ID  \
    --file=IWOSaAccessOrg.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 Intersight Workload Optimizer to execute actions in your GCP environment.

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

```
title: "IWO Role: Project Action Execution"
           description: "Grant IWO 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 `IWOSaProjectAction.yaml` with the path to your yaml file:

```
$ gcloud iam roles create $CUSTOM_ROLE_NAME \
    --project=$PROJECT_ID  \
    --file=IWOSaProjectAction.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 Intersight Workload Optimizer.

   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 Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer

Create a Custom Billing Role

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

```
title: "IWO Billing Data Viewer Role"
description: "Minimal Required Permissions for IWO to view billed cost and pricing stored in the GCP Pro
ject"
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 `IWOSaBilling.yaml` with the path to your yaml file:

```
$ gcloud iam roles create $CUSTOM_ROLE_NAME \
    --project=$PROJECT_ID  \
    --file=IWOSaBilling.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 Intersight Workload Optimizer.

   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

Intersight Workload Optimizer 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 154)*
- EnablingWinRMViaPowerShell *(on page 155)*

Enabling WinRM Via Global Policy Objects

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

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

1. On the AD domain controller, open the Group Policy Management Console (GPMC). If the GPMC is not installed, see https://technet.microsoft.com/en-us/library/cc725932.aspx.
2. Create a new Global Policy Object:
   a. In the GPMC tree, right-click **Group Policy Objects** within the domain containing your Hyper-V servers.
   b. Choose **Create a GPO in this domain,** and link it here.
   c. Enter a name for the new GPO and click **OK**.
3. Specify the computers that need access:
   a. Select the new GPO from the tree.
   b. On the **Scope** tab, under **Security Filtering**, specify the computer or group of computers you want to grant access. Make sure you include all of your Hyper-V targets.

4. Right-click the new GPO and choose **Edit** to open the Group Policy Management Editor.
5. Configure the WinRM Service:
   a. In the Group Policy Management Editor, select **Computer Configuration > Policies > Administrative Templates > Windows Components > Windows Remote Management (WinRM) > WinRM Service**.
   b. Double-click each of following settings and configure as specified:

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

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

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

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

Allow Remote Shell Access:	Enabled

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

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

```
gpupdate /force
```

Enabling WinRM Via PowerShell

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

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

   ```
   hp-vx485
   hp-vx486
   ```

2. Since Cisco requires changes to the default quickconfig settings, create a batch file containing the following command:

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

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

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

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

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

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

Port Configuration

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

NOTE:
This list may include targets not available to your version or distribution of Intersight Workload Optimizer.

Target	**Ports**
VMWare vCenter (Monitoring)	443
VMWare vCenter (Tags)	10443
VMWare vCenter (Kubernetes)	88
Microsoft Hyper-V	5985, 5986
Microsoft Hyper-V (Kubernetes)	88
Microsoft VMM	5985, 5986
CloudFoundry	80, 443
Dell VMAX	5988, 5989
Nutanix	9440
PureStorage	80, 443
HPE 3PAR	5988, 5989, 8080
NetApp	80, 443
Cisco UCS	80, 443
IBM WebSphere	8880
Oracle WebLogic	7001
Apache Tomcat	1009
Oracle Database	1521
Microsoft SQL Server	1433
MySQL	3306
IBM FlashSystem	7443
HPE OneView	
AWS; Microsoft Azure; Google Cloud Platform	
kubeturbo	
AppDynamics	

Target	Ports
Azure AppInsights	
DynaTrace	
Flexera	
Instana	
NewRelic	
ServiceNOW	

AWS Target IAM Role Requirements

This appendix guides you through the steps to configure the Intersight Workload Optimizer 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.

NOTE:
The Intersight Workload Optimizer SaaS offering also supports adding AWS targets via IAM roles. See https://support-turbonomic.force.com/TurbonomicCustomerCommunity/s/article/Turbonomic-SaaS-IAM-Role-Setup.

Pre-requisites

Follow the steps below to configure the AWS Mediation pods to leverage IAM roles:

1. Follow the instructions provided by Amazon EKS, 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
   - 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer default Service Account with which the AWS Mediation pods are running.

NOTE:
Unless specified in the Intersight Workload Optimizer 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.
   - Intersight Workload Optimizer 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 Intersight Workload Optimizer 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 Intersight Workload Optimizer AWS Account Target setup by going to the Intersight Workload Optimizer UI to configure the AWS Account targets using an IAM role. See Amazon Web Services *(on page 8)*.