Cisco HyperFlex Sizer Getting Started Guide

First Published: 2017-05-16

Last Modified: 2021-01-06

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Communications, Services, Bias-free Language, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Bias-Free Language

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

Communications, Services, Bias-free Language, and Additional Information

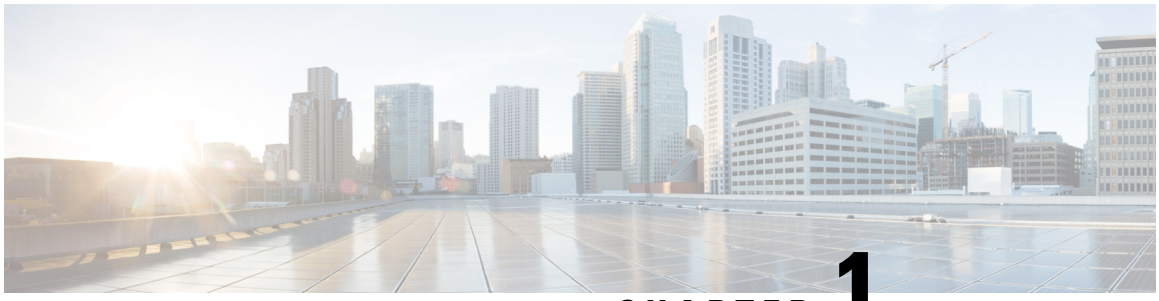

CHAPTER 1

Overview

Introduction

Cisco HyperFlex Sizer is a web-based application that helps in sizing different workloads and matching them with appropriate Cisco Hyperconverged Systems.

The HyperFlex Sizer supports the following workloads:

- Virtual Desktop Infrastructure (VDI)
- General Server Virtualized Environment (VSI)
- Microsoft SQL database
- Oracle
- Microsoft Exchange Server
- Compute and Capacity Sizer (RAW)
- HX Edge (ROBO)

Installation Prerequisites

Supported Browsers

Browser	Version
Chrome	65 or higher
Firefox	59 or higher

Access HyperFlex Sizer

HyperFlex Sizer is hosted on the Cisco Lightweight Application Environment (LAE) infrastructure. You can access the HyperFlex Sizer using the following link:

https://hyperflexsizer.cloudapps.cisco.com

Enter your Cisco user credentials.

User Name	Cisco user ID
Password	Cisco password

Click **Log In**.

HyperFlex Sizer Home Page

The HyperFlex Sizer home page displays a list of all the scenarios previously created by you, if any. This is the first page you see when you log in to HyperFlex Sizer.

UI Element	Description
Add Scenario button	If there are no scenarios, you can click the **Add Scenario** button to create a new scenario. If there are any previously created scenarios, the **Add Scenario** button is on the top-right corner.
My Scenario tab	Displays a list of all scenarios previously created by you, if any.
Shared with me tab	Displays a list of all the shared scenarios by other users, if any.
Sizing Calculator button	Provides a link to the sizing calculator, which is a tool that calculates the effective resources available after reserve and overhead reductions for a given node configuration.
Download HX Tools button	Provides links to download the OVA files for HxBench and HxProfiler.
Getting Started button	Provides training materials for HxSizer, HxBench, and HxProfiler.
What's New button	Provides information on what's new in the various HyperFlex Sizer releases.
Send Feedback button	Click the Envelope icon to send feedback and/or queries, if any.

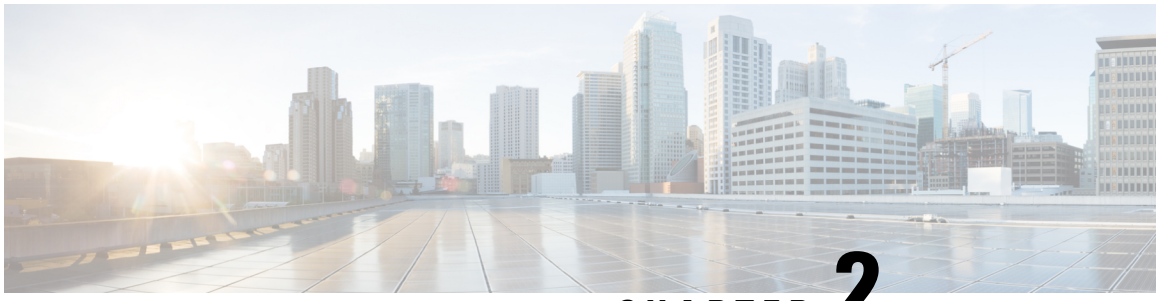

CHAPTER 2

Scenarios

Add a Scenario

The following steps describe how to add a Scenario:

Step 1 In the HyperFlex Sizer home page, click the **Add Scenario** button. The Add Scenario window appears, shown as follows:

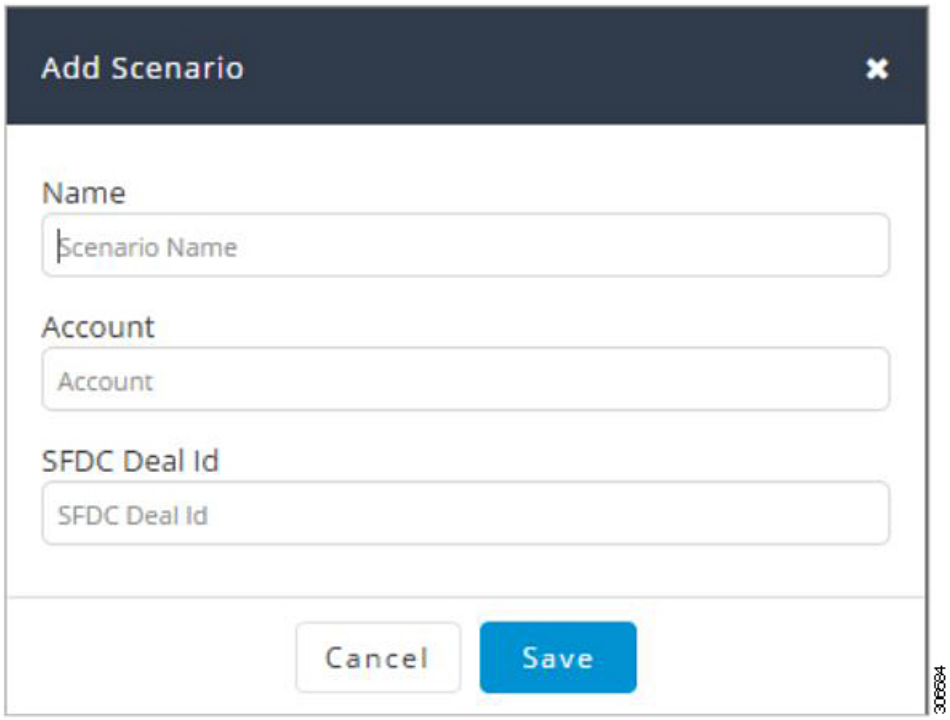

Step 2 In the **Add Scenario** window, complete the following fields:

Field Name	Description
Name	Enter a name for the sizing Scenario. Use the following guidelines to create the name: • The Scenario name must begin with an alphabetic character. • Use alphanumeric characters only. An underscore, hyphen, or plus symbol can be used as a separator. • Special characters are not allowed. • The Scenario name must be unique.
Account	(Optional)
SFDC Deal Id	(Optional)

Step 3 Click **Save**.

You are now redirected to the Scenario details page.

Scenarios Page with Workloads

The HyperFlex Sizer Scenario details page displays a list of all the Workloads created by you, if any.

You can size different workloads on the Scenario details page by using the various options provided on the HyperFlex Sizer web application. You can choose between All- Flash and Lowest Cost sizing options to view the recommended sizing configurations that you can use in your HyperFlex cluster.

The node and parts for sizing are chosen based on CPU normalization.

Normalized cores:

The performance of a core in one processor is different from that of another processor. The performance of a CPU varies across generations of processors of the same type. The HyperFlex Sizer computes the effective core for a node by using the SpecInt and CFP values, and normalizes this to the SpecInt/CFP values of Intel E5-2630 v4.

Resize button—The Resize button helps to resize the Scenario to the latest version in case of a version mismatch for older Scenarios.

Note Cisco recommends using the **All-Flash** option as it gives high performance and optimal sizing options for your HyperFlex cluster.

The fields described in this section are displayed under both the **All-Flash** tab and the **Lowest_Cost** tab, shown as follows:

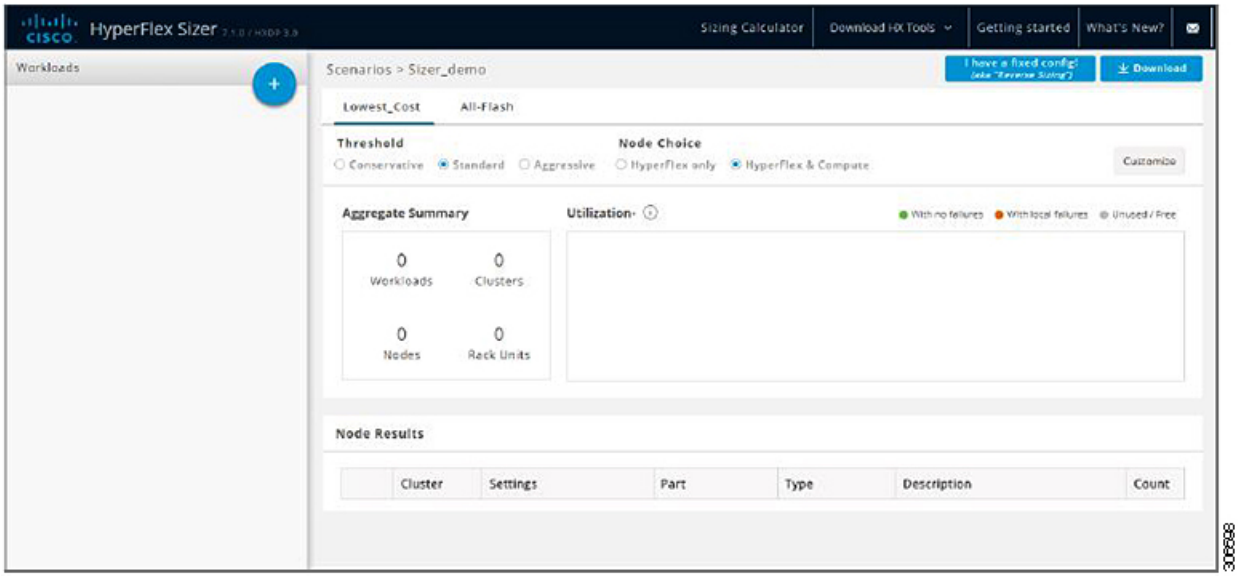

Scenario Details Page

UI Element	Description
Threshold check box	Set the sizing threshold to one of the following: • **Standard**—Default • **Conservative** • **Aggressive**
Node Choice check box	Choose the type of node for which you want to calculate sizing: • **HyperFlex & Compute**—Default • **HyperFlex Only**
Aggregate Summary column	Displays the recommended number of Clusters, Nodes, and Rack Units for a given number of workloads.
Utilization column	Displays the expected hardware resource utilization, namely CPU, RAM, Storage Capacity, and Storage IOPS for a workload. There are three different components under utilization: • **With no failures**—For workloads that have replication enabled, the resource utilization includes the replication overhead. • **With local failures**—The number of failures matches the sizing parameter used for specifying the *Performance Headroom* parameter. For workloads that have replication enabled, resource utilization includes the replication overhead. • **Unused / Free**—Applicable only if the workload has replication enabled. Shows the resource utilization when the DR partner cluster has failed and workloads running on the cluster have moved over.
Node Results column	You can view the combined Node Results for all clusters or for the individual clusters. The following results are displayed: • **Cluster**—Name of the HX cluster • **Settings**—Specific settings for the HX cluster, for instance, the Replication Factor that is set • **Part**—Part number • **Type**—Type of node • **Description**—Properties of the node • **Count**—Number of nodes

Customize Button

Click **Customize** to custom select the type for HX cluster workload for sizing and click **Apply**. You can create customized sizing for the workload using the following options:

UI Element	Description
Threshold button	Set the sizing threshold to one of the following: • **Standard**—Default • **Conservative** • **Aggressive**
Include Software Cost button	• **Yes** • **No**
Node Choice button	Choose the type of node for which you want to size: • **HyperFlex & Compute**—Default • **HyperFlex Only**
Sizing Option button	• **Bundle Only**—Size the bundle only. • **Bundle & CTO**—Size both bundle and Configure to Order bundle.
Server Type button	Choose the type of server for which you want to size: • **All** • **M4 Server** • **M5 Server**
Disk Option button	Select the type of disk: • **All** • **SED Only** (Self-Encrypting Drive Only) • **NVMe** (Non-Volatile Memory Express)
Modular LAN button	• **All** • **40G VIC**
You can customize **HyperFlex Nodes**, **Compute Nodes**, **CPU**, **RAM Slots**, **Disk Options**, and **Modular LANs**.	

Download Button

You can download the sizing reports from the HxSizer in two formats:

Download Sizing Report

You can view all details of the sizing input, proposed sizing configurations, workload summary, aggregate workload requirements, and utilization of resources for the chosen option. Click the Download button in the top-right corner of the Scenario page (shown as follows), then click Download Sizing Report.

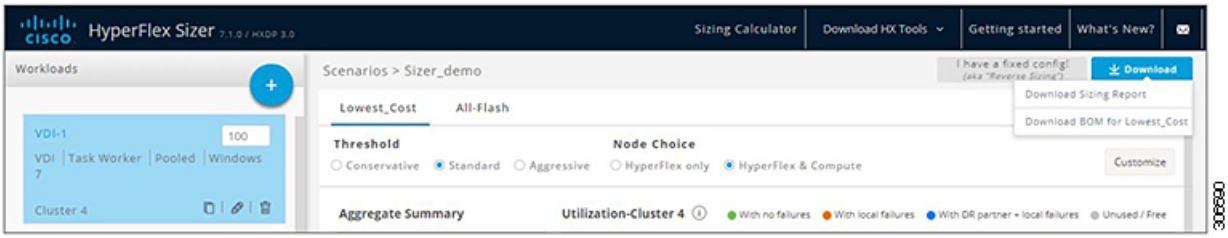

Download BOM for All-Flash and Lowest_Cost

A detailed Bill of Materials (BOM) is available separately for the All-Flash and Lowest Cost options as an Excel spreadsheet. This Excel sheet can be directly loaded to your Cisco Commerce Workspace (CCW).

Fixed Configuration (Reverse Sizing)

In the Fixed Configuration (also referred to as "Reverse Sizing"), the workflow starts with a fixed HX configuration and helps validate whether a given set of workloads will run on it or not. Whereas in Regular Sizing, the workflow helps identify the cost-optimal HX configuration for a set of workloads. See Fixed (Reverse) Configuration Sizing, on page 36 for more information.

Scenario Tasks

To view the existing Scenarios, navigate to the **My Scenarios** tab in HyperFlex Sizer. You can perform the following tasks with existing Scenarios.

Clone a Scenario

Click the *Clone* icon of an existing Scenario to create a copy of the Scenario and complete the following fields.

Field Name	Description
Name	Enter a name for the sizing Scenario. Use the following guidelines to create the name: • The Scenario name must begin with an alphabetic character. • Use alphanumeric characters only. An underscore, hyphen, or plus symbol can be used as a separator. • Special characters are not allowed. • The Scenario name must be unique.
Account	(Optional)

Field Name	Description
SFDC Deal Id	(Optional)

Edit a Scenario

Click the *Edit* icon of an existing Scenario to edit the **Scenario name**, **Account**, or **SFDC Deal ID**.

Share a Scenario

To share a Scenario with an existing user or a new user, follow these steps:

1. Click the *Share* icon of an existing Scenario to share the Scenario.
2. Click the **Select User** drop-down list, to add user(s) from the populated list.
3. You can set the following access privileges to users:
   - Read-only access—Users can only access the Scenario, they cannot modify it.
   - Read and write access—Users have privileges to modify the Scenario, add new workloads, and modify existing workloads.

   If username is not available in the populated list, provide a valid User ID under the **User Name** field and press **Enter**.
4. Click **Save**.

You can find the list of Scenarios shared with you under the **Shared with me** tab. You can find details about the owner of the Scenario and the users with whom the Scenario is shared on the **Shared Scenarios** page.

Delete a Scenario

Click the *Delete* icon of an existing Scenario to delete the Scenario.

Modify VM/Desktop Counts

To change the default values, click **Customize**.

Attention The recommended values are based on performance tests and should be changed with caution.

To modify the number of desktops/VMs:

Step 1 Change the value in the Workloads List for any workload in the Count box (shown as follows).

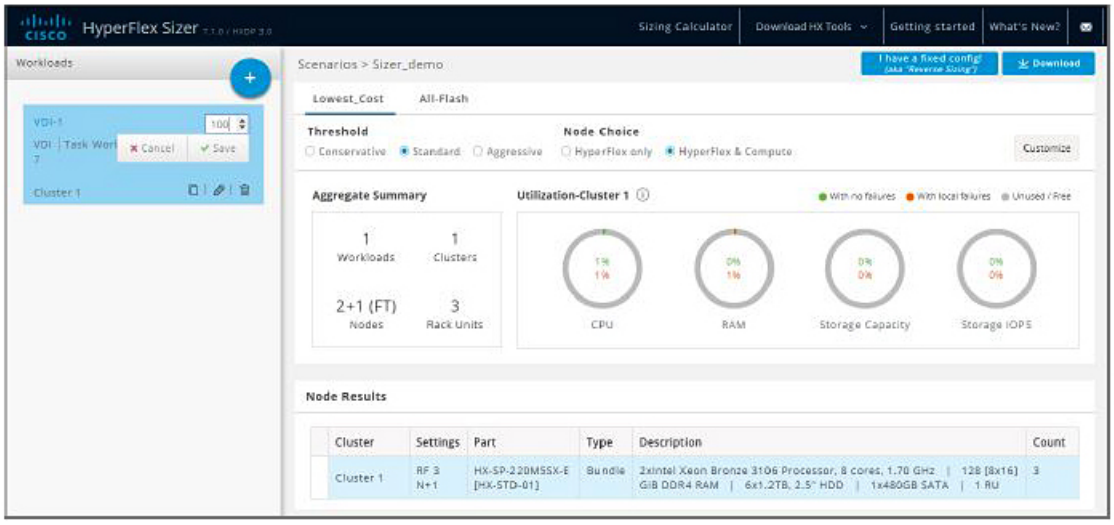

Step 2 Click **Save**.

View Clusterwise Workload Results

To view clusterwise Workload results :

Step 1 Click Cluster 1 under Node Results to see all the Workloads that belong to Cluster 1. The Summary and Utilization details are displayed based on Cluster 1.

Step 2 Click **Cluster 2** under Node Results to view the Cluster 2 details.

Customize Options for Sizing

To customize options for sizing:

Step 1 Using the Customize button on the top-right screen, customize the Hyperflex Nodes, Compute Nodes, CPU, RAM Slots, and RAM options (shown as follows):

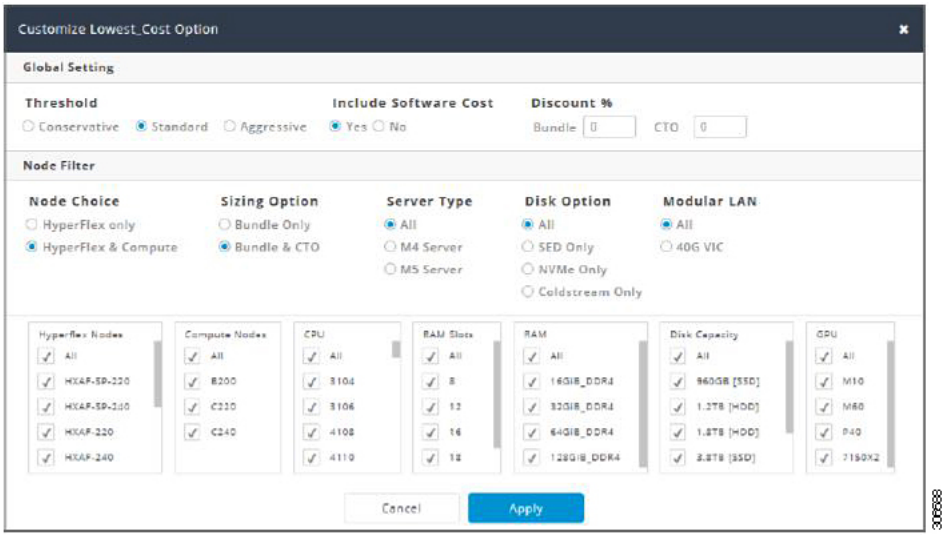

Step 2 You can also choose to customize the Threshold, Node Choice, Sizing Option, and Disk Option. The Disk Option allows you to select Storage Encrypted Disks, NVMe, or Coldstream.

Step 3 Click **Apply**. The changed options are saved and the new result can be viewed from the Scenario Page.

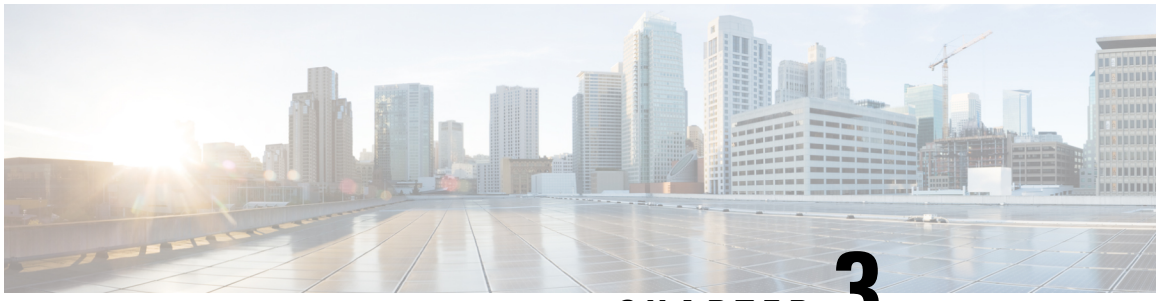

CHAPTER 3

Workloads

Workloads Pane

You can complete the following actions from the Workloads Pane:

Clone Workload

Click the *Clone* icon of an existing Workload to create a copy of the Workload and complete the following fields.

Edit Workload

Click the *Edit* icon of an existing Workload to edit the Workload profile.

Delete Workload

Click the *Delete* icon of an existing Workload to delete the Workload.

Modify Virtual Machine or Desktop Count

To modify the number of desktops or VMs for a Workload, change the value in the **Count** box. Click **Save**.

Add VDI Workload

To change the default values, click **Customize**.

Attention The recommended values are based on performance tests and should be changed with care.

To add a VDI Workload:

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** page, select **VDI**, (shown as follows). Click **Next**.

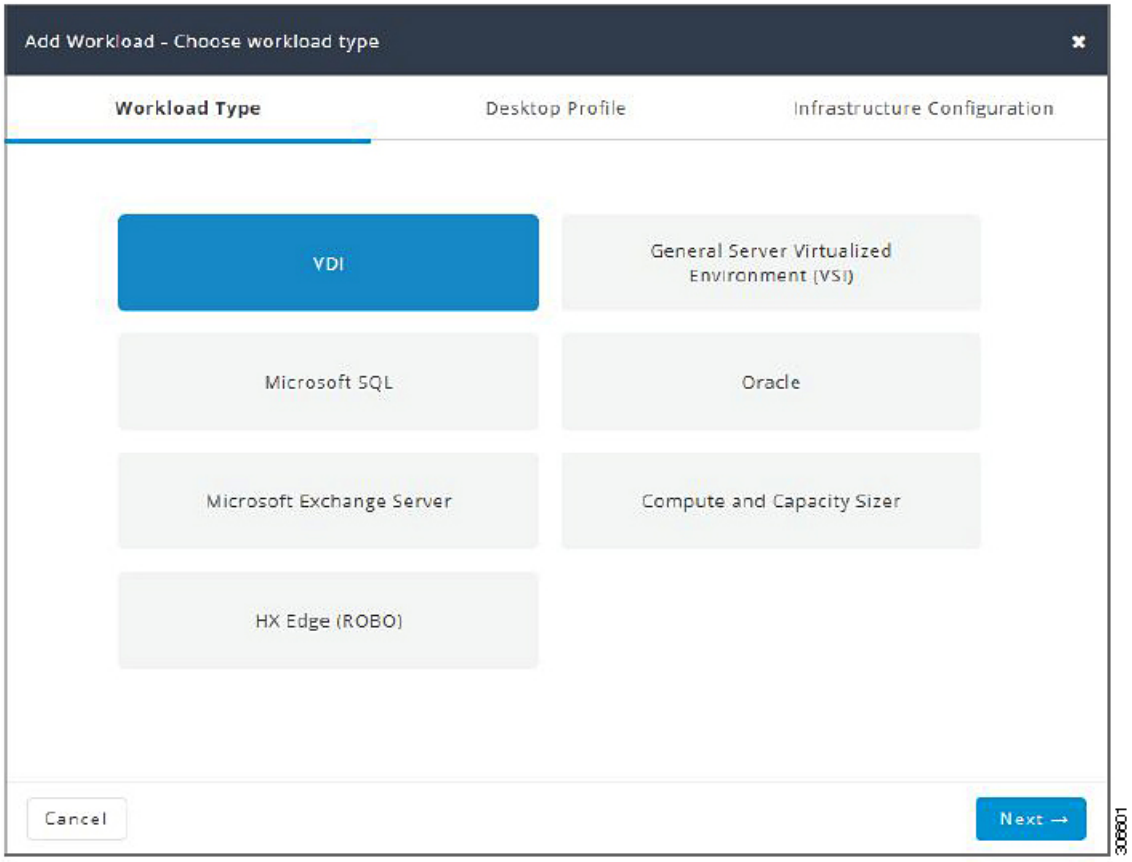

Step 3 On the **Desktop Profile** page (shown as follows), complete the following fields:

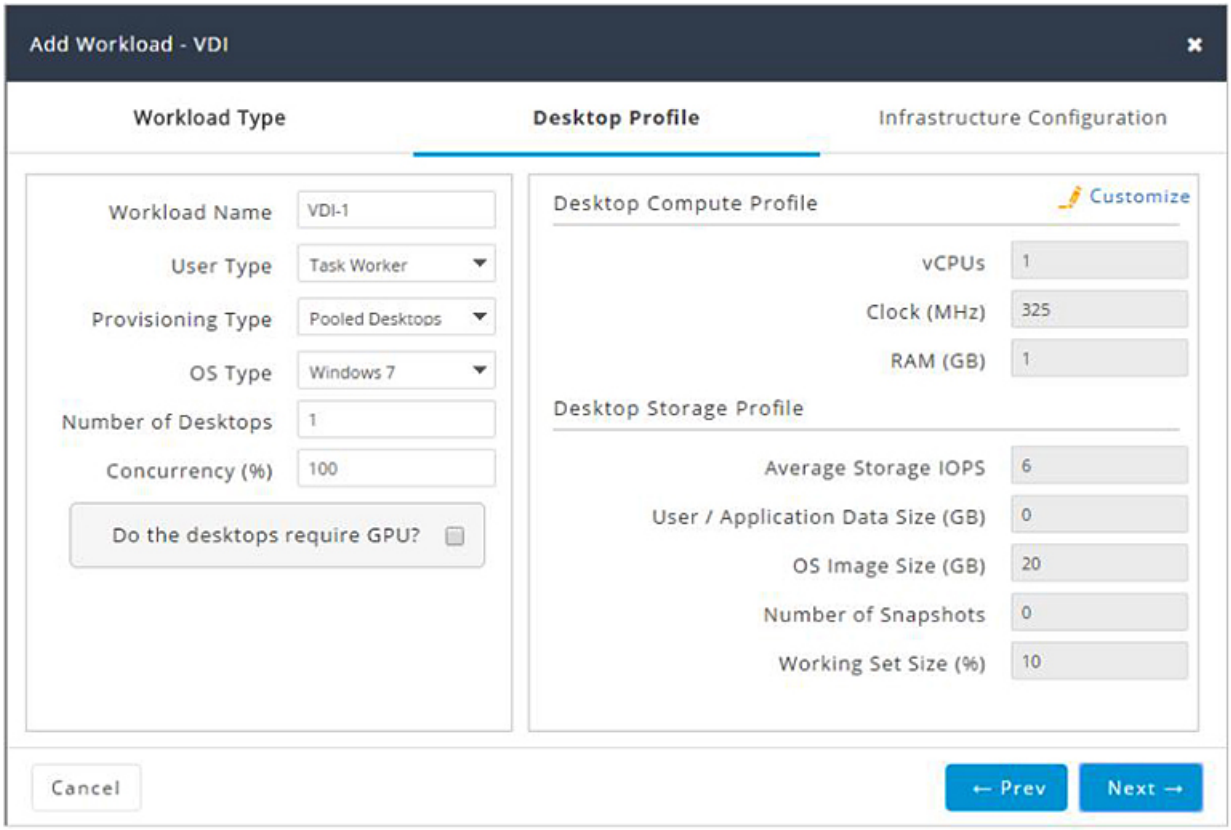

UI Element	Description
Workload Name field	Name of the Workload
User Type drop-down list	Choose from a list of predefined resource consumption values: • **Task Worker** • **Knowledge Worker** • **Power User** • **Custom User**—If the predefined resource consumption values in the templates listed do not meet your requirements, select the Custom User option to manually enter the Desktop Compute Profile and Desktop Storage Profile values.
Provisioning Type drop-down list	You have the following options for data retention: • **Persistent Desktops**—Retains data on the desktop. • **Pooled Desktops**—Does not retain on the desktop.

UI Element	Description
OS Type drop-down list	• **Windows 7** • **Windows 10**
Number of Desktops field	Enter the total number of desktops. The limit is 1 - 30,000 desktops.
Concurrency (%) field	Enter percentage relevant to the total number of desktops that should remain powered on concurrently.
Do the desktops require GPU?	Indicate if the desktops need to use GPUs.
Desktop Compute Profile Depending on the User Type you choose, the recommended values will change.	
vCPUs field	• Task Worker—1 vCPU • Knowledge Worker—2 VCPUs • Power User—2 VCPUs
Clock (MHz) field	• Task Worker—325 MHz • Knowledge Worker—400 MHz • Power User—400 MHz
RAM (GB) field	• Task Worker—1 GB • Knowledge Worker—2 GB • Power User—2 GB
Desktop Storage Profile	
Average Storage IOPS field	Depending on the User Type you choose, the recommended values will change. • Task Worker—6 IOPs • Knowledge Worker—8 IOPs • Power User—10 IOPs
User / Application Data Size (GB) field	Recommended is 0 GB
OS Image Size (GB) field	Recommended is 20 GB
Number of Snapshots field	Recommended is 0 GB
Working Set Size (%) field	Recommended is 10%

Click **Next**.

Step 4 On the **Infrastructure Configuration** page (shown as follows), complete the following fields.

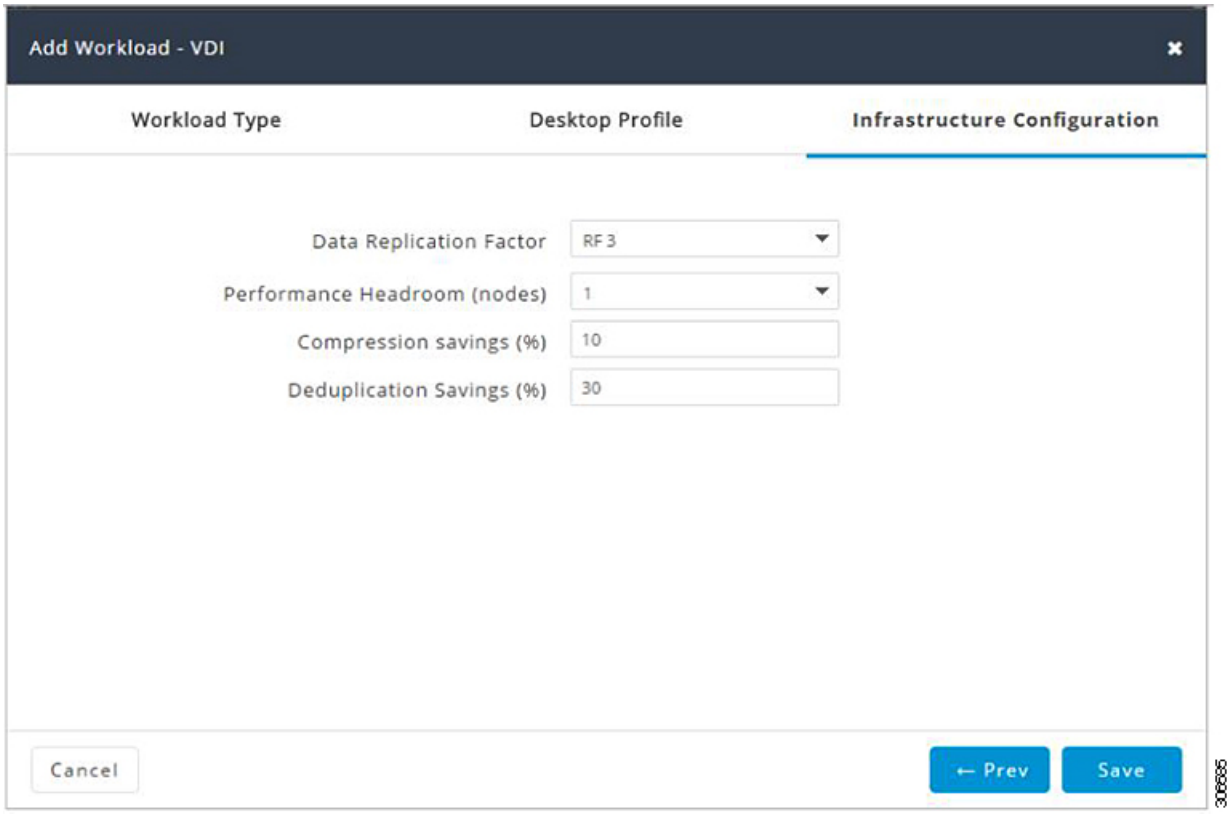

UI Element	Description
Data Replication Factor drop-down list	RF3 is recommended for data redundancy.
Performance Headroom (nodes) drop-down list	Enter the number of nodes used for Fault Tolerance. Recommended is 1 node. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of a node failure.
Compression Savings (%) field	Recommended is 10%
Deduplication Settings (%) field	Recommended is 30%

Step 5 Click **Save**.

Add General Server VSI Workload

To change the default values, click **Customize**.

Attention The recommended values are based on performance tests and should be changed with care.

To add a General Server Virtualized Environment (VSI) Workload:

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** page, select **General Server Virtualized Environment (VSI)** (shown as follows). Click **Next**.

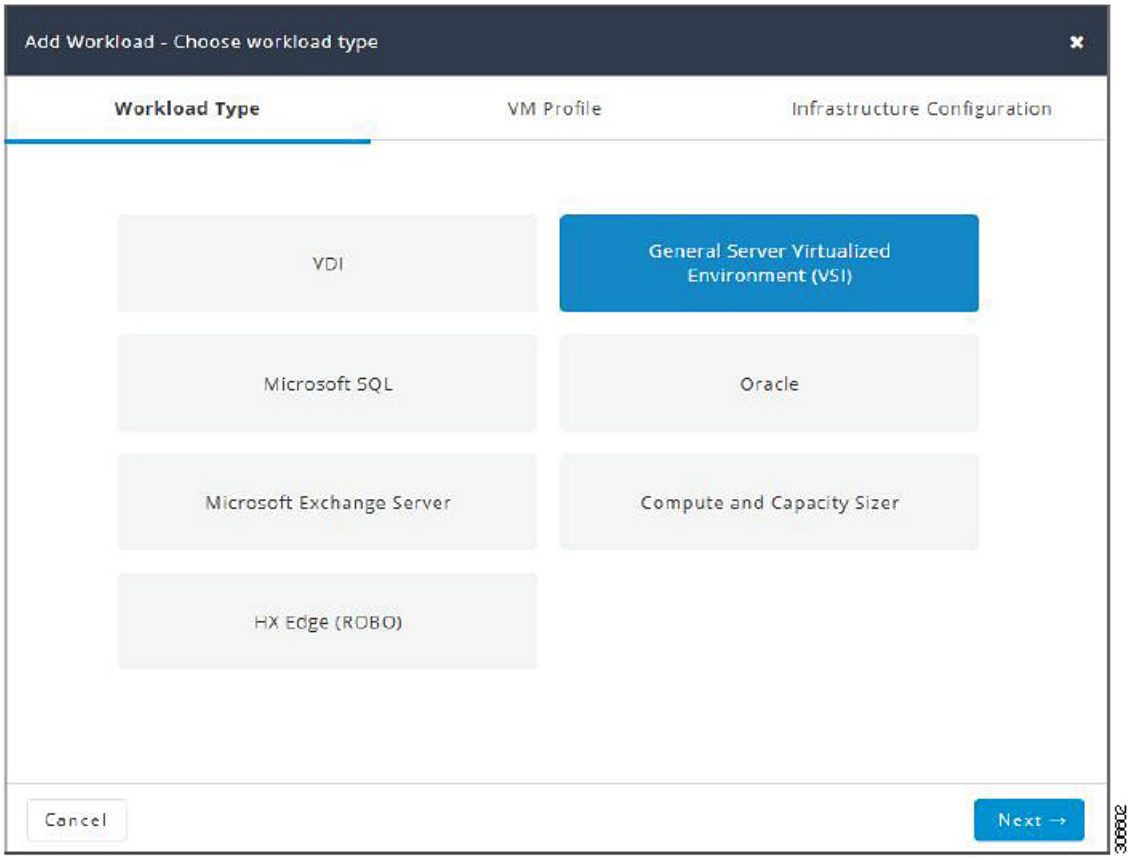

Step 3 On the **VM Profile** page, complete the following fields:

UI Element	Description
Workload Name field	Enter a name for the Workload.

UI Element	Description
VM Type drop-down list	Choose from a list of predefined resource consumptions values: • **Small** • **Medium** • **Large** • **Custom**—If the predefined resource consumption values in the templates listed do not meet the requirements, select Custom option to enter profile values on the Infrastructure Configuration page.
Number of VMs field	Enter the number of VMs.
VM Compute Profile Depending on the VM Type you choose, the recommended values will change.	
vCPUs field	• Small—2 vCPUs • Medium—4 vCPUs • Large—8 vCPUs
vCPU Overprovisioning Ratio field	Recommended value for all VM Types is 4 vCPUs. The total number of vCPUs that can be packed per core.
RAM (GB) field	• Small—8 GB • Medium—16 GB • Large—32 GB
VM Storage Profile Depending on the VM Type you choose, the recommended values will change.	
Average 8K Storage IOPS field	• Small—50 IOPS • Medium—100 IOPS • Large—200 IOPS
User / Application Data Size (GB) field	• Small—50 GB • Medium—200 GB • Large—750 GB
OS Image Size (GB) field	Recommended is 20 GB. Size of the OS image for the VM.

UI Element	Description
Number of Snapshots field	Recommended is 5 snapshots.
Working Set Size (%) field	Recommended is 10%

Click **Next**.

Step 4 On the **Infrastructure Configuration** page, complete the following fields.

UI Element	Description
Cluster Type button	• Normal • Stretch - The Stretch Cluster provides a high-availability cluster for data of high importance. This cluster is spread across two geographic regions and will be available even if one site goes down completely for any reason, such as a natural disaster.
Data Replication Factor drop-down list	RF2 is recommended for better availability.
Performance Headroom (nodes) drop-down list	Enter the number of nodes used for Fault Tolerance. Recommended is 1 node. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of a node failure.
Compression Savings (%) field	Recommended is 20%
Deduplication Savings (%) field	Recommended is 10%
Enable Remote Replication?	Choose to enable remote replication. You can now set Workload placement and site failure protection as follows: **Primary Workload Placement** drop-down list • **Site A** • **Site B** **Site Failure Protection (% Workload)**—Recommended is 100.

Step 5 Click **Save**.

Add Microsoft SQL Workload

To change the default values, click **Customize**.

Attention The recommended values are based on performance tests and should be changed with care.

To add a Microsoft SQL Workload:

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** page, select **Microsoft SQL** (shown as follows). Click **Next**.

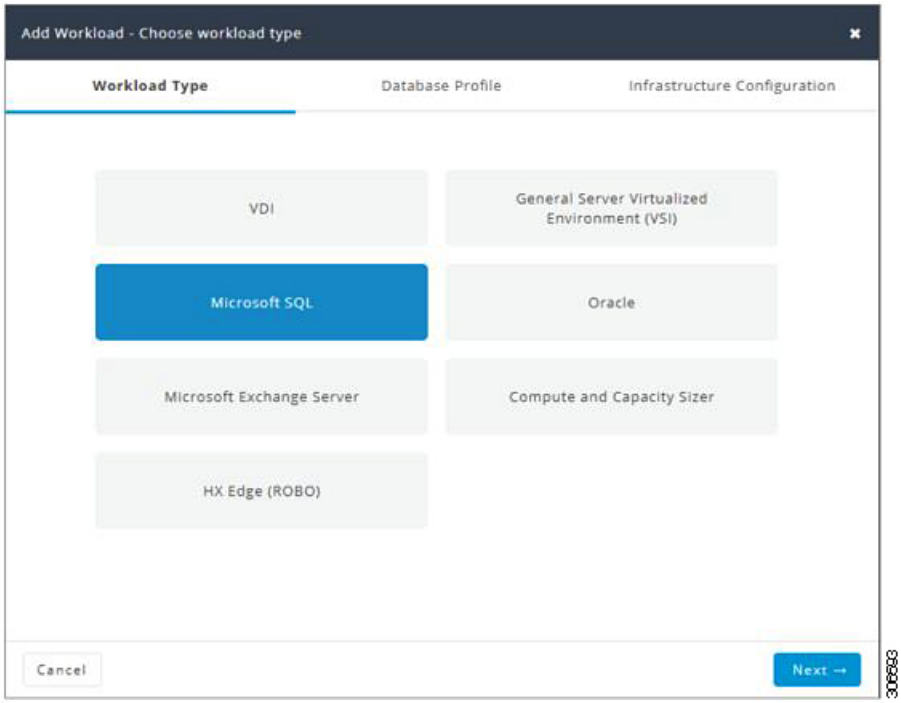

Step 3 On the **Database Profile** page, complete the following fields:

UI Element	Description
Workload Name field	Enter a name of the Workload.
Database Type drop-down list	You can choose OLTP or OLAP database type. • **OLTP**—Represents transactional workloads. The Sizer assigns a workload that consists of 8K 70% read, 30% write; 100% random, when sizing for the specified number of IOPS for OLTP. • **OLAP**—Represents query, reporting, or analytics workloads. The Sizer assigns a workload that consists of large sequential reads when sizing for the specified throughput for OLAP.

UI Element	Description
Database Profile drop-down list	Choose from a list of predefined Database Profile values: • **Small** • **Medium** • **Large** • **Custom**—If the predefined values in the templates listed do not meet your requirements, select the **Custom** option to manually enter Compute Profile and Storage Profile values.
Number of Databases field	Enter the total number of databases.
Compute Profile Depending on the Database Profile you choose, the recommended values will change.	
vCPUs field	• Small—2 vCPUs • Medium—4 vCPUs • Large—8 vCPUs
vCPU Provisioning Ratio field	Recommended is 2 vCPUs.
RAM (GB) field	• Small—8 GB • Medium—16 GB • Large—32 GB
Storage Profile Depending on the Database Profile you choose, the recommended values will change.	
Database Size (GB) field	• Small—400 GB • Medium—1000 GB • Large—4000 GB

UI Element	Description
IOPS field	IOPS changes based on the Database Type you choose. For OLTP Database Type, the following values are recommended: • Small—1000 IOPS • Medium—3000 IOPS • Large—10000 IOPS For OLAP Database Type, the following values are recommended: • Small—100 MB/s • Medium—200 MB/s • Large—800 MB/s
Database Overhead (%) field	• Small—45% • Medium—40% • Large—30%

Click **Next**.

Step 4 On the **Infrastructure Configuration** page, complete the following fields:

UI Element	Description
Cluster Type button	• Normal • Stretch - The Stretch Cluster provides a high-availability cluster for data of high importance. This cluster is spread across two geographic regions and will be available even if one site goes down completely for any reason, such as a natural disaster.
Data Replication Factor drop-down list	RF3 is recommended for data redundancy.
Performance Headroom (nodes) drop-down list	Enter the number of nodes used for Fault Tolerance. Recommended is 1 node. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of a node failure.
Compression Savings (%) field	Recommended is 30%
Deduplication Settings (%) field	Recommended is 0%

UI Element	Description
Enable Remote Replication? check box	Choose to enable remote replication. You can now set Workload placement and site failure protection as follows: **Primary Workload Placement** drop-down list • **Site A** • **Site B** **Site Failure Protection (% Workload)**—Recommended is 100%

Step 5 Click **Save**.

Add Oracle Workload

To change the default values, click **Customize**.

Attention The recommended values are based on performance tests and should be changed with caution.

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** page, select **Oracle** (shown as follows). Click **Next**.

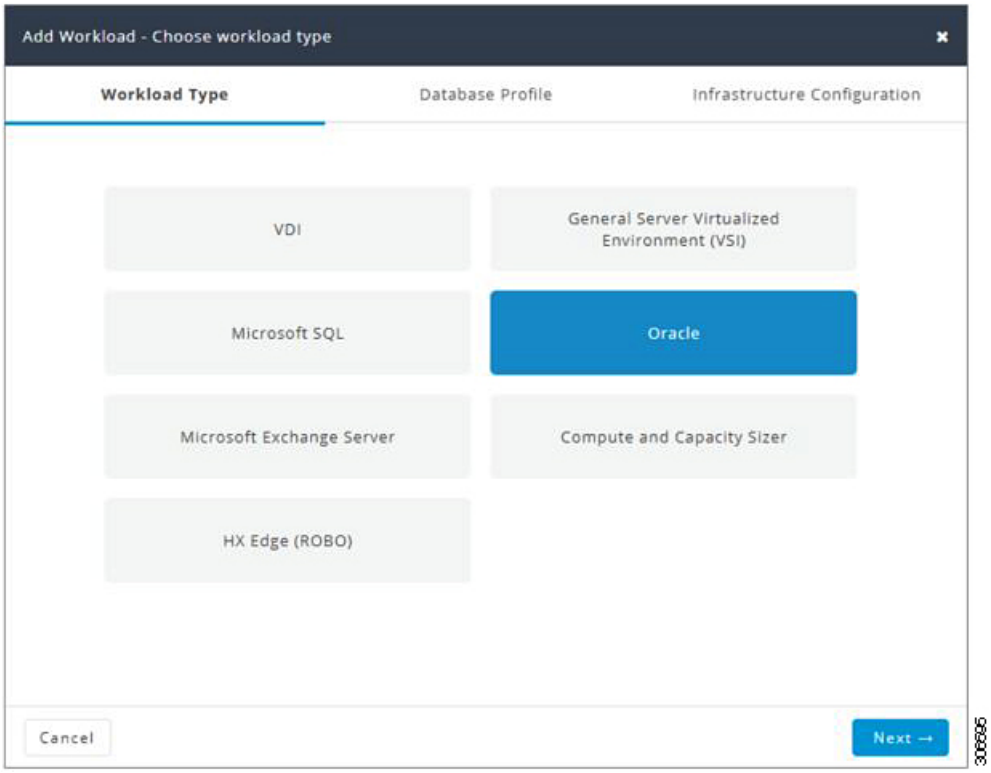

Step 3 On the **Database Profile** page, complete the following fields:

UI Element	Description
Workload Name field	Enter a name of the Workload.
Database Type drop-down list	You can choose OLTP or OLAP database type. • **OLTP**—Represents transactional workloads. The Sizer assigns a Workload that consists of 8K 70% read, 30% write; 100% random, when sizing for the specified number of IOPS for OLTP. • **OLAP**—Represents query, reporting, or analytics workloads. Sizer assigns a workload that consists of large sequential reads, when sizing for the specified throughput for OLAP.

UI Element	Description
Database Profile drop-down list	Choose from a list of predefined Database Profile values: • **Small** • **Medium** • **Large** • **Custom**—If the predefined values in the templates listed do not meet your requirements, select the Custom option to manually enter Compute Profile and Storage Profile values.
Number of Databases field	Enter the total number of databases.
Compute Profile Depending on the Database Profile you choose, the recommended values will change.	
vCPUs field	• Small—4 vCPUs • Medium—8 vCPUs • Large—16 vCPUs
vCPU Provisioning Ratio field	Recommended is 2 vCPUs.
RAM (GB) field	• Small—16 GB • Medium—64 GB • Large—96 GB
Storage Profile Depending on the Database Profile you choose, the recommended values will change.	
Database Size (GB) field	• Small—400 GB • Medium—1000 GB • Large—4000 GB

UI Element	Description
IOPS field	IOPS changes based on the Database Type you choose. For OLTP Database Type, the following values are recommended: • Small—6000 IOPS • Medium—10000 IOPS • Large—30000 IOPS For OLAP Database Type, the following values are recommended: • Small—200 MB/s • Medium—400 MB/s • Large—1000 MB/s
Database Overhead (%) field	• Small—45% • Medium—40% • Large—30%

Click **Next**.

Step 4 On the **Infrastructure Configuration** page, complete the following fields:

UI Element	Description
Cluster Type button	• Normal • Stretch - The Stretch Cluster provides a high-availability cluster for data of high importance. This cluster is spread across two geographic regions and will be available even if one site goes down completely for any reason, such as a natural disaster.
Data Replication Factor drop-down list	RF3 is recommended for data redundancy.
Performance Headroom (nodes) drop-down list	Enter the number of nodes used for Fault Tolerance. Recommended is 1 node. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of a node failure.
Compression Savings (%) field	Recommended is 30%
Deduplication Settings (%) field	Recommended is 0%

UI Element	Description
Enable Remote Replication? check box	Choose to enable remote replication. You can now set Workload placement, and site failure protection as follows: **Primary Workload Placement** drop-down list • **Site A** • **Site B** **Site Failure Protection (% Workload)**—Recommended is 100.

Step 5 Click **Save**.

Add Microsoft Exchange Server Workload

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** page, select **Microsoft Exchange Server** (shown as follows). Click **Next**.

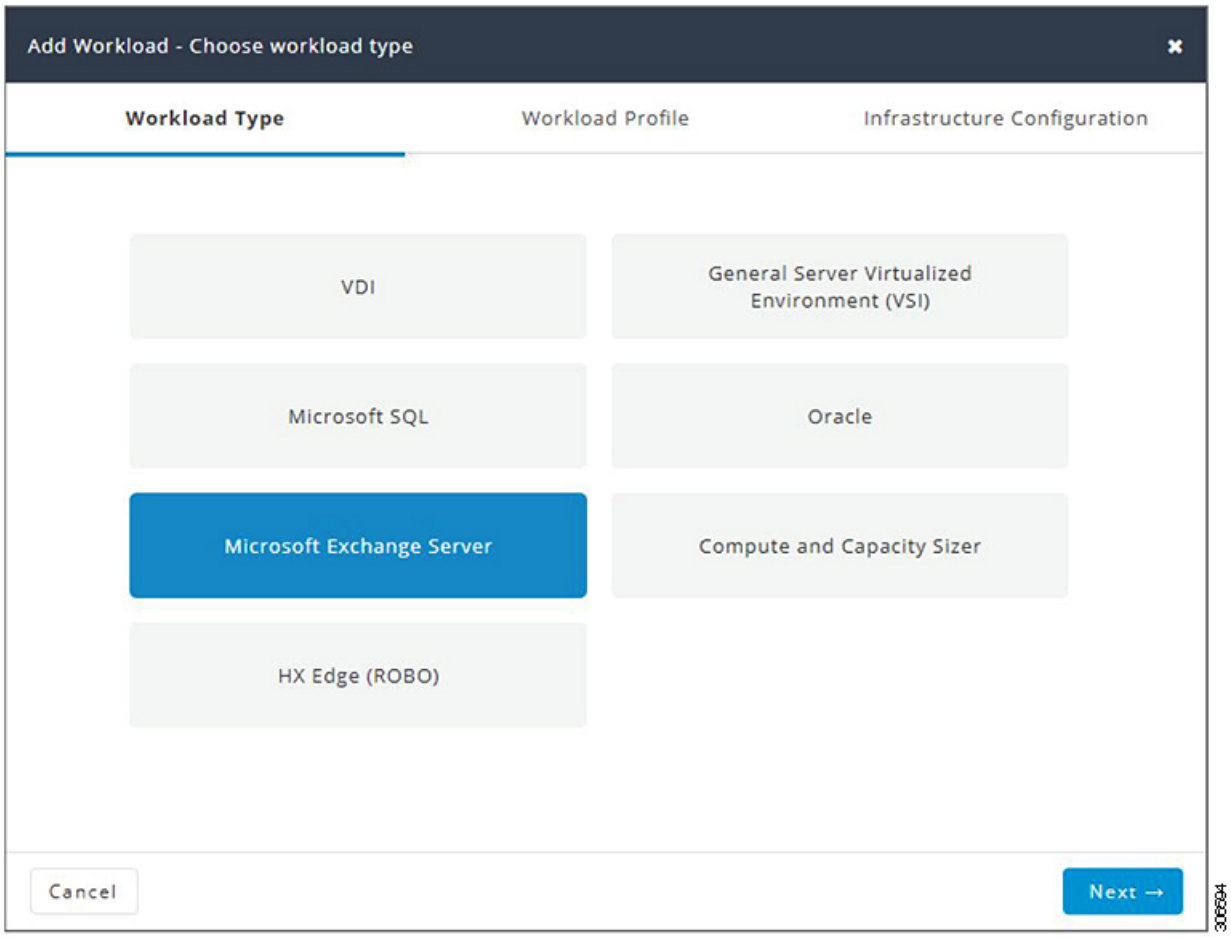

Step 3 On the **Workload Profile** page, you can choose to enter the values manually or you can import them from a file.

UI Element	Essential Information
Workload Name field	Enter a name for the Workload.
Workload Input Type	Download the Microsoft Exchange Workload modeling spreadsheet from Microsoft Exchange 2013 Server Role Requirements Calculator. **Important** Ensure that the Microsoft Exchange 2013 Sizing Calculator is filled out properly, refer to the Configure the Microsoft Exchange 2013 Server Role Requirements Calculator, on page 41. Upload the completed *.XLSM* spreadsheet to process workload inputs.
vCPUs field	Total number of cores required for all the MS Exchange Servers after accounting for system overhead. Intel E5-2630 v4 is used as the reference CPU for core count.

UI Element	Essential Information
vCPU Overprovisioning Ratio field	Total number of vCPUs that can be packed per core.
Total RAM (GB) field	The total RAM required for all guest VMs, after accounting for system overhead.
Effective User Capacity (GB) field	This value depends on the Dedupe or Compression savings. You can change the Deduplication and compression savings on the Infrastructure Configuration Page.
DB IOPS field	Average 16KB IOPS, with 100% random 60/40 read/write ratio.
Log IOPS field	Average 32KB IOPS, with 100% random 60/40 read/write ratio.
Maintenance IOPS field	Average 64KB IOPS, with 100% random 60/40 read/write ratio.
Future Growth (%) field	Specify percentage to allow for future growth of the environment for Physical Cores, RAM, and Effective User Capacity.

Click **Next**.

Step 4 On the **Infrastructure Configuration** page, complete the following fields.

UI Element	Essential Information
Cluster Type button	• Normal • Stretch - The Stretch Cluster provides a high-availability cluster for data of high importance. This cluster is spread across two geographic regions and will be available even if one site goes down completely for any reason, such as a natural disaster.
Data Replication Factor field	RF3 is recommended for better availability.
Performance Headroom (# of nodes) field	Number of nodes of Fault Tolerance. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of node failure.
Compression Savings (%) field	By default is set to 15%. The allowed range is 0-50%
Deduplication Settings (%) field	By default is set to 15%. The allowed range is 0-70%

Step 5 Click **Save**.

Add Compute and Capacity Sizer (RAW) Workloads

To add the Compute and Capacity Sizer Workloads:

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** tab, select **Compute and Capacity Sizer** (shown as follows). Click **Next**.

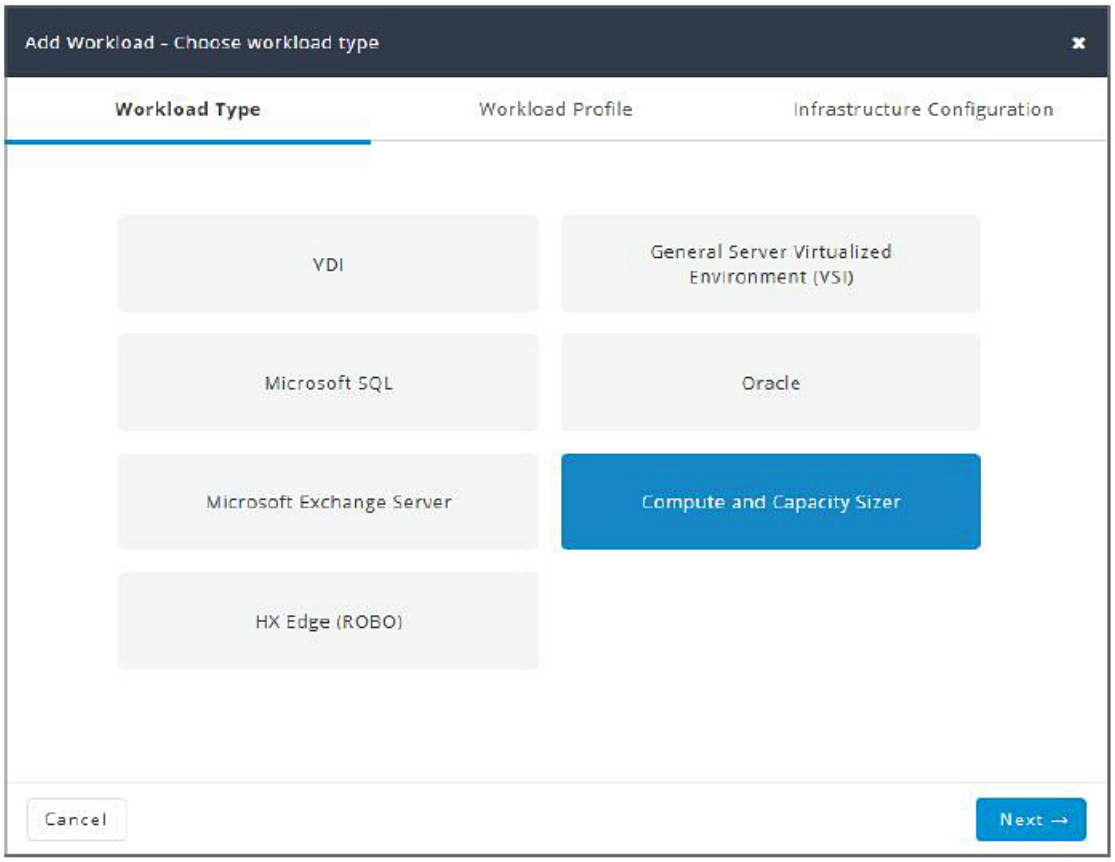

Step 3 On the **Workload Profile** page, you can choose to enter the values manually or you can import them from a file.

UI Element	Description
Workload Name field	Enter a name for the Workload.

UI Element	Description
Workload Input Type button	• **Manual**—To use default values, choose this option. • **File**—You can import the values from a CSV file. The CSV file can be downloaded from the HxProfiler. The File option provides users with the following options: • 30-day summary from the HX Profiler tool. (The CSV file can be downloaded from the HX Profiler tool for a 30-day duration period.) • RV Tools Output
CPU Unit field	• Cores by default • Clock
Total vCPUs field	Default is 2 vCPUs. The total number of cores required for all the guest VMs after accounting for system overhead.
CPU Overprovisioning Ratio field	Default is 1 vCPU. The total number of vCPUs that can be packed per core.
Total RAM (GB) field	Default is 128 GB. The total RAM required for all guest VMs after accounting for system overhead.
Effective User Capacity (GB) field	Default is 1000 GB. This value depends on the dedupe or compression savings. You can change the deduplication and compression savings on the Infrastructure Configuration page.
Future Growth (%) field	Specify the percentage to allow for future growth of the environment for Physical Cores, RAM, and Effective User Capacity.

Click **Next**.

Step 4 On the **Infrastructure Configuration** page, complete the following fields.

UI Element	Description
Cluster Type button	• Normal • Stretch - The Stretch Cluster provides a high-availability cluster for data of high importance. This cluster is spread across two geographic regions and will be available even if one site goes down completely for any reason, such as a natural disaster.
Data Replication Factor field	RF3 is recommended for better availability.
Performance Headroom (# of nodes) field	Number of nodes of Fault Tolerance. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of a node failure.
Compression Savings (%) field	By default set to 0%. The allowed range is 0-50%.
Deduplication Settings (%) field	By default set to 0%. The allowed range is 0-70%.

Step 5 Click **Save**.

Add HyperFlex Edge (ROBO) Workload

To change the default values, click **Customize**.

Attention The recommended values are based on performance tests and should be changed with care.

To add a HyperFlex Edge (ROBO) Workload:

Step 1 Click the + icon under **Workloads**.

Step 2 On the **Workload Type** page, select **HX Edge (ROBO)** (shown as follows). Click **Next**.

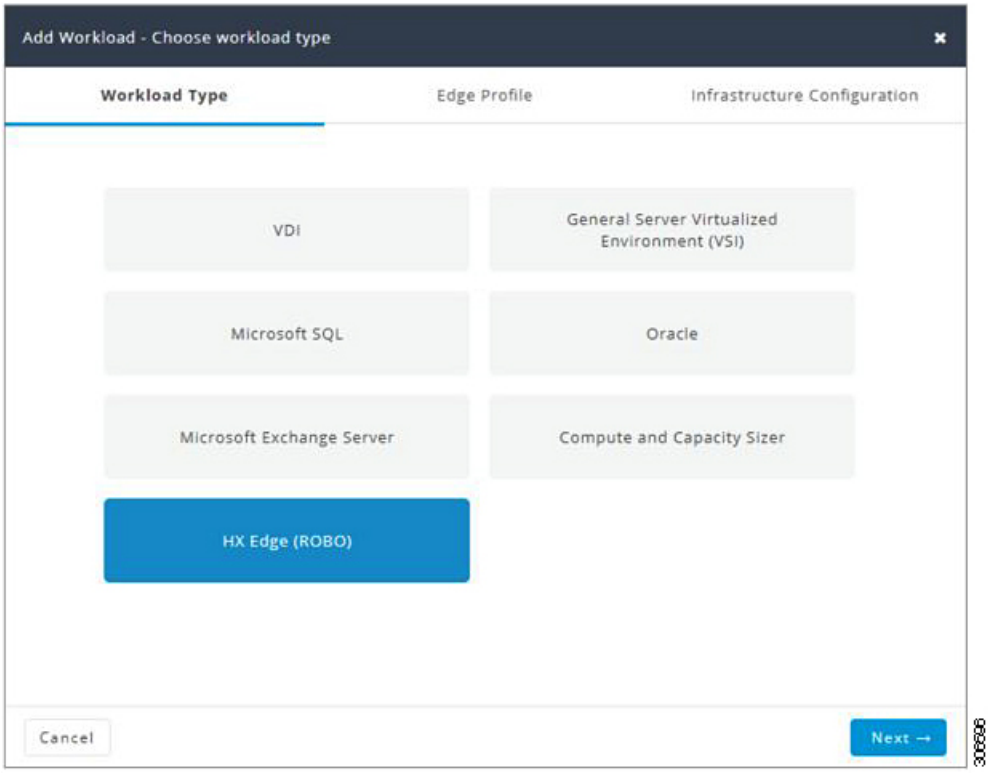

Step 3 On the **Edge Profile** page, complete the following fields:

UI Element	Description
Workload Name field	Enter a name for the Workload.
VM Type drop-down list	Choose from a list of predefined resource consumptions values: • **Small** • **Medium** • **Large** • **Custom**—If the predefined resource consumption values in the templates listed do not meet the requirements, select the Custom option to enter profile values on the Infrastructure Configuration page.
Number of VMs field	Enter the number of VMs.
VM Compute Profile Depending on the VM Type you choose, the recommended values will change.	

UI Element	Description
vCPUs field	• Small—2 vCPUs • Medium—4 vCPUs • Large—8 vCPUs
vCPU Overprovisioning Ratio field	Recommended value for all VM Types is 4. The total number of vCPUs that can be packed per core.
RAM (GB) field	• Small—8 GB • Medium—16 GB • Large—32 GB
VM Storage Profile Depending on the VM Type you choose, the recommended values will change.	
Average 8K Storage IOPS field	• Small—50 IOPS • Medium—100 IOPS • Large—200 IOPS
User / Application Data Size (GB) field	• Small—50 GB • Medium—100 GB • Large—750 GB
OS Image Size (GB) field	Recommended is 20 GB. Size of the OS image for the VM.
Number of Snapshots field	Recommended is 5 snapshots
Working Set Size (%) field	Recommended is 10%

Click **Next**.

Step 4 On the **Infrastructure Configuration** page, complete the following fields.

UI Element	Description
Data Replication Factor drop-down list	**Caution** Edge workload is supported only with RF 2.
Performance Headroom (nodes) drop-down list	Enter the number of nodes used for Fault Tolerance. Recommended is 1 node. Setting Performance Headroom adds additional nodes to the cluster to ensure that there is enough performance bandwidth in case of a node failure.

UI Element	Description
Compression Savings (%) field	Recommended is 20%
Deduplication Savings (%) field	Recommended is 10%

Step 5 Click **Save**.

Fixed (Reverse) Configuration Sizing

Fixed Sizing (also referred to as "Reverse Sizing") is a workflow that starts with a fixed configuration and helps validate whether a given set of Workloads will run on it or not. In the case of Regular Sizing, the workflow helps identify the cost-optimal HX configuration for a set of Workloads.

To add a Fixed Configuration Sizing workflow:

Step 1 Create a new Scenario, then click I have fixed config (aka "Reverse Sizing"). Click Yes to confirm (shown as follows).

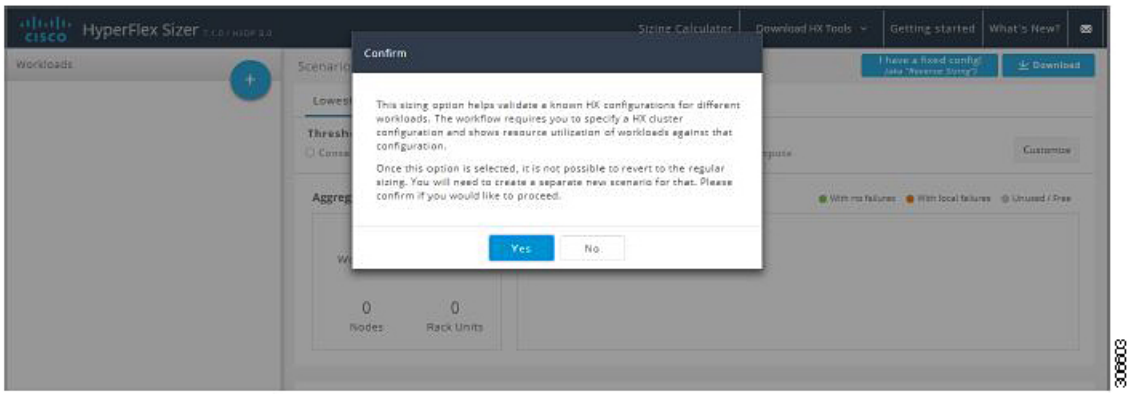

Note Once you have changed to a fixed configuration (Fixed Sizing), you cannot revert back to Regular Sizing.

Step 2 The Fixed Config Sizing tab appears with options to select the HyperFlex node and the Compute node (shown as follows). Make your selections, then click Apply. The Scenario page reloads.

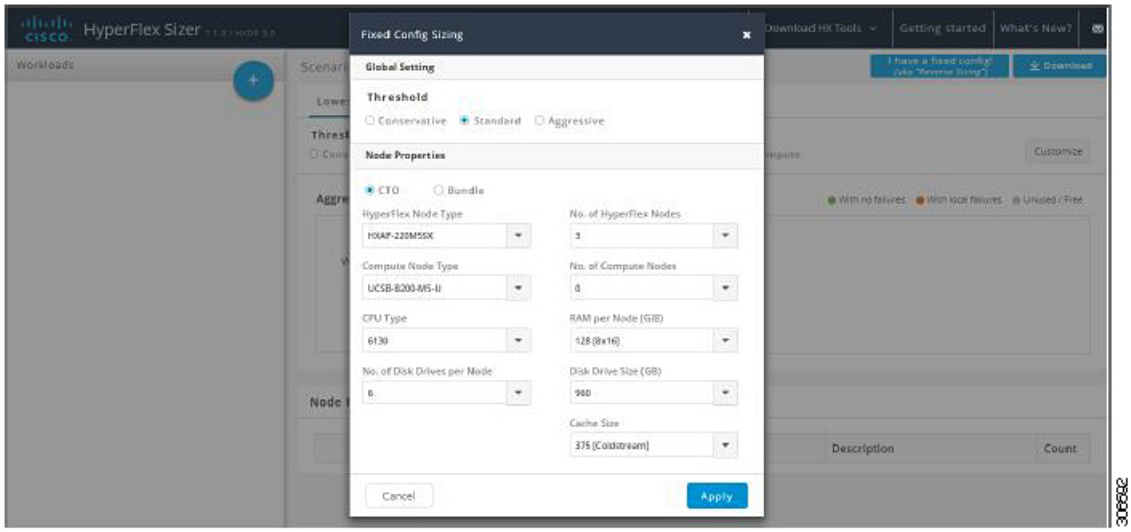

Step 3 Click the + icon under **Workloads**, which prompts a dialogue box with the various Workload types supported (shown as follows). Any greyed-out Workloads are not supported for the chosed fixed configuration. This setting can be changed by selecting the Customize button on the Scenario page.

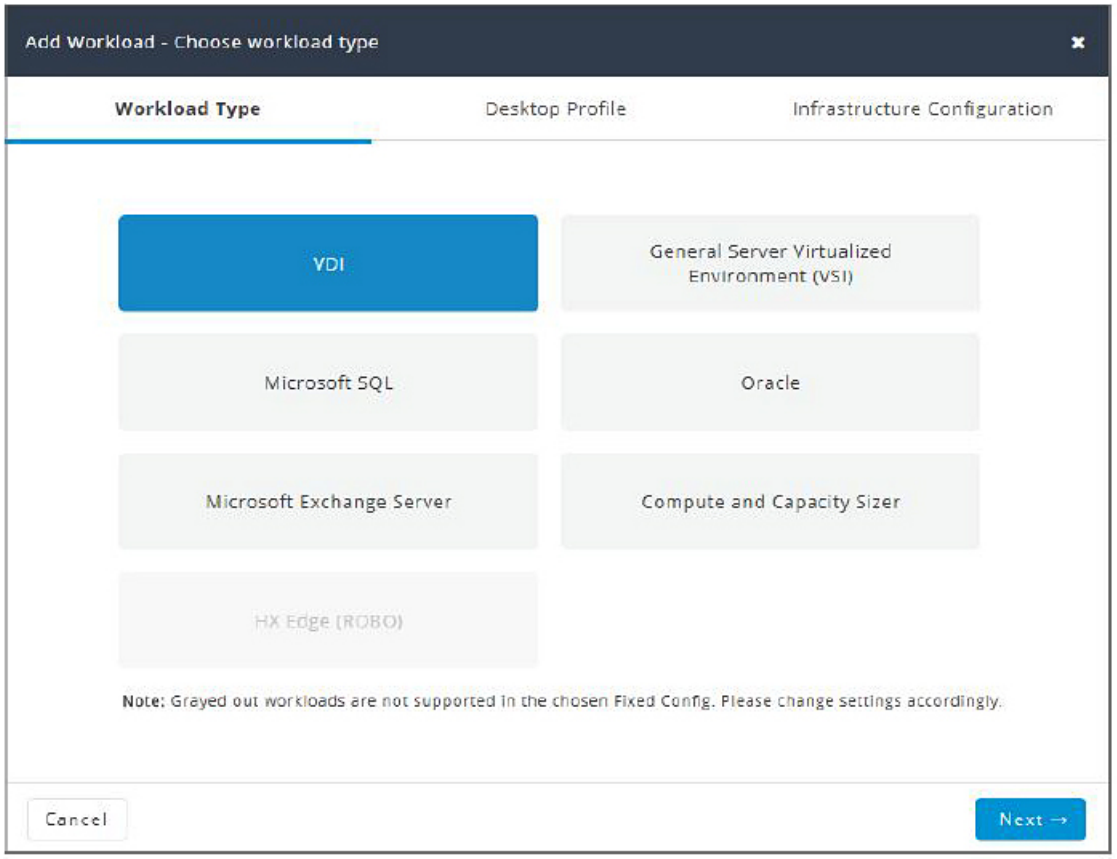

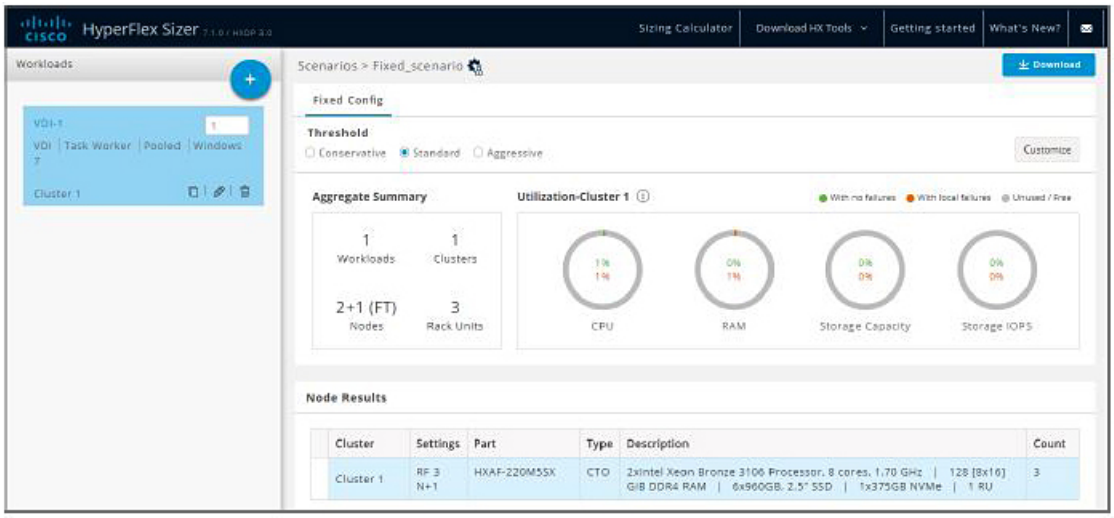

The other Workloads can be added to the Fixed Configuration Sizing based on the clustering of those Workloads that can be placed into one cluster. The standard clustering formats include:

`[VDI], [VSI, DB, ORACLE], [RAW], [EXCHANGE], [ROBO]`

Note For Fixed Configuration Sizing, the Stretch Cluster and Replication are not supported.

Sizing Calculator

The Sizing Calculator is a tool that calculates the effective resources available after reserve and overhead reductions for a given node configuration.

To use the Sizing Calculator:

Step 1 Select the required configurations and click the Calculate button, shown as follows.

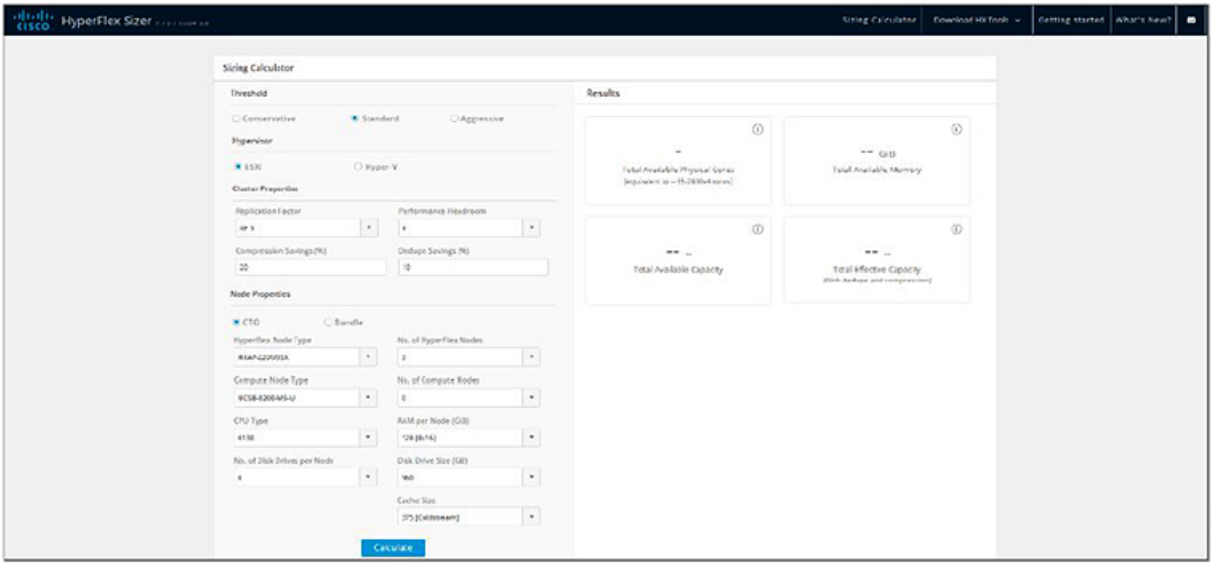

Step 2 Click Download Report (shown as follows) if you want to download the Sizing Calculator report.

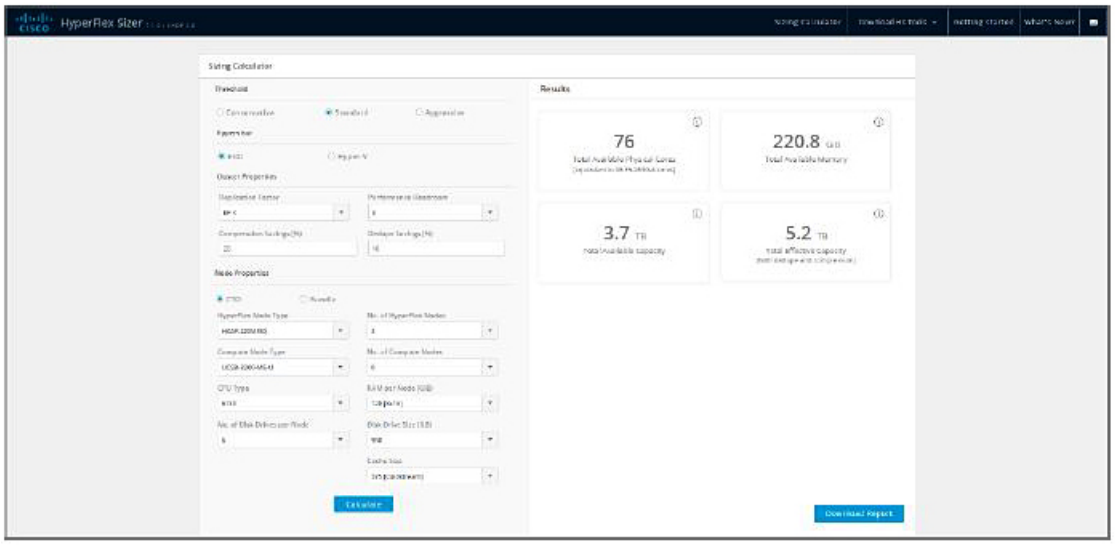

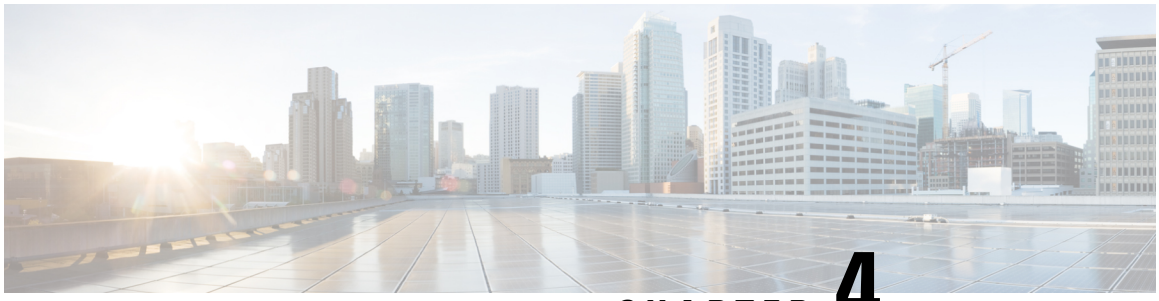

CHAPTER 4

Configure Microsoft Exchange 2013 Server Role Requirements Calculator

Configure the Microsoft Exchange 2013 Server Role Requirements Calculator

Overview

Download the Microsoft Exchange Workload modeling spreadsheet from Microsoft Exchange 2013 Server Role Requirements Calculator. Read the Microsoft Exchange Calculator Readme file for comprehensive guidance on using the calculator.

Cisco HyperFlex Sizer provides the BOM for the primary datacenter only. This section provides the parameters that should be configured on the **Input Tab** of the Microsoft Exchange Calculator. Customers who plan to deploy in multiple datacenters and stretch the Database Availability Group (DAG) must complete the input for the secondary datacenter under **Site Resilience Configuration**. Completing this input ensures that the primary datacenter Compute and Storage requirements are properly sized to handle all users in the event that the secondary datacenter is down.

Exchange Environment Configuration

Configuration Settings	Value
Exchange Server Version	2016
Global Catalog Server Architecture	64-bit
Server Role Virtualization	Yes
High Availability Deployment	Yes If a DAG is planned, ensure that *High Availability Deployment* is set to **Yes**and that the proper number of database copy instances are selected for each site.

Tier-1 [2,3,4] User Mailbox Configuration

Ensure that the user mailbox tiers are set to the appropriate initial and maximum mailbox sizes. With HyperFlex, adding additional persistent tier disks or adding converged nodes to the cluster expands the usable storage on the cluster automatically. Adding additional databases, expanding the HyperFlex datastore, or expanding the Windows LUN where an online database is located, is instant and can occur without any downtime.

Configuration Settings	Value
User Mailbox Configuration Settings	
Number of Days in a Work Week field	5 days
Tier-1 User Mailbox Configuration field	
Total Number of Tier-1 User Mailboxes / Environment field	10000 Tier-1 User Mailboxes/Environment
Projected Mailbox Number Growth Percentage field	0%
Total Send/Receive Capability / Mailbox / Day field	200 messages
Average Message Size (KB) field	75 KB
Initial Mailbox Size (MB) field	2048 MB
Mailbox Size Limit (MB) field	10240 MB

Backup Configuration

Configuration Settings	Value
Backup Methodology field	Backup Methodology can have a sizing impact. The recommended methodology is to utilize the Cisco HyperFlex native snapshot, along with a third-party backup application that will keep a copy of the backup off the cluster, typically in a backup repository. The following backup methodology options are available: • (Recommended) Hardware VSS Backup/Restore—Requires the smallest amount of capacity to restore LUNs. • Software VSS Backup/Restore—Requires a larger capacity to restore LUNs. • Exchange Native Data Protection—Requires a larger capacity to restore LUNs. • VMware redo-log snapshots—Large restore LUNs must be provisioned on each Exchange Server to provide enough space to pull a copy of the backup and then enable roll forward recovery. • Exchange Native Data Protection with lagged database copies—Large restore LUNs must be provisioned on each Exchange Server to provide enough space to pull a copy of the backup and then enable roll forward recovery.
Backup Frequency field	Weekly Full or Daily Incremental
Backup/Truncation Failure Tolerance field	3
Network Failure Tolerance (Days) field	0 days

Storage Options

Configuration Settings	Value
Automatically Calculate Number of Exchange Database Volumes Required	Yes If set to *No*, carefully size and ensure that enough Exchange Data Volumes are selected so that the databases fit on the server. If they do not fit, you will be unable to upload the spreadsheet into the Cisco HyperFlex Sizer, and a warning will appear on cell G216 on the Role Requirements tab.
Number of AutoReseed Volumes per Server	1 AutoReseed Volumes

Server Configuration

The Microsoft Exchange Calculator is based on a particular baseline CPU. To properly calculate the megacycle to actual CPU consumption, enter both the number of vCPUs for the Microsoft Exchange Server VM in *Processor Cores / Server*, and the *SPECint2006 Rate Value* for the Cisco HyperFlex Server. For example values, refer to SPEC CINT2006 Result.

Server Configuration	Processor Cores / Server	SPECint2006 Rate Value
Primary Datacenter Mailbox Servers	16	2330
Secondary Datacenter Mailbox Servers	16	2330

Troubleshooting

Error Message	Recommended Solution
One or more workloads have exceeded the maximum CPU limits.	Toggle to *HX+Compute*, or include *All-Flash* options if not already set. Split the Workload into smaller Workloads.
No SmartPlay hyperconverged nodes have been chosen, due to filters. Please change the filters.	For *All-Flash* Option, select *SmartPlay Hyperconverged Nodes* from the *Customize* option.