



Cisco HyperFlex Workload Profiler for Kubernetes

First Published: 2022-02-11

Last Modified: 2022-03-24

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

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

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

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

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

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

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

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

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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2022 Cisco Systems, Inc. All rights reserved.



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.



CHAPTER 1

Introduction

- [Introduction, on page 1](#)

Introduction

Kubernetes Introduction

HX Workload Profiler is a Kubernetes characterization tool that can estimate the compute, storage, and network usage of hosts and containers across multiple Kubernetes deployments. The output of the tool can be fed into HXSizer product to size the HyperFlex cluster requirements. The workload requirements are estimated by the HX Profiler tool.

HX Workload Profiler key outputs:

- 30-day historic summary of host-level compute metrics
- Detailed compute, storage, and network metrics for hosts and Containers for all Kubernetes profiled
- Detailed time series data to analyze trends of key metrics in the UI
- Download of time series data of all key metrics

The HX Workload Profiler tool is delivered as a tar file and can be used to monitor multiple Kubernetes servers simultaneously.



CHAPTER 2

System Requirements

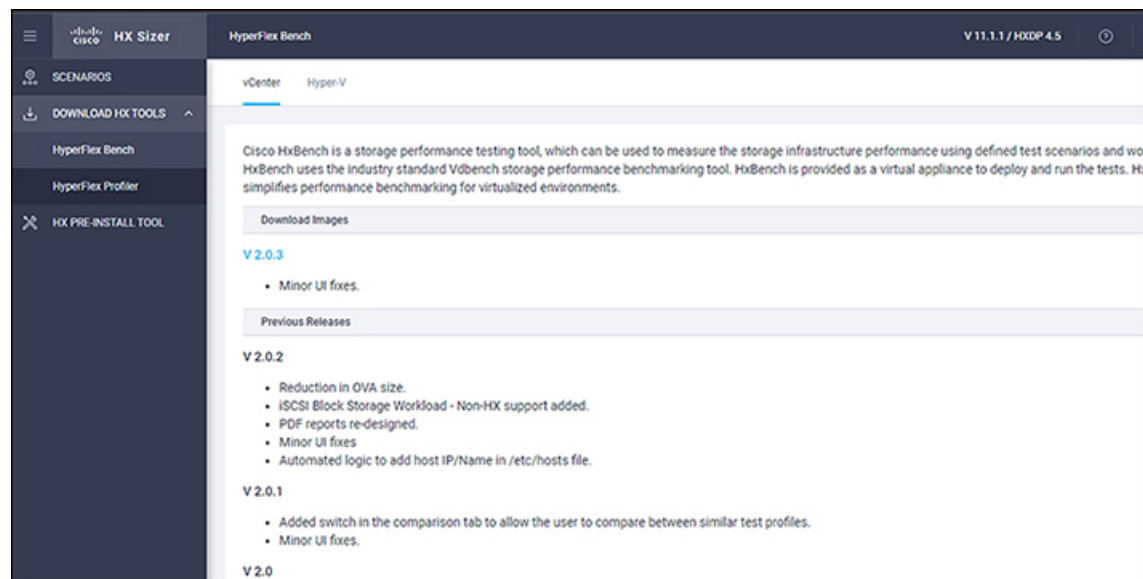
- [Prerequisites, on page 3](#)
- [Configuring HX Profiler for Kubernetes Environments, on page 3](#)
- [Using the Profiler Service, on page 8](#)
- [Locating the Kubernetes Application Logs, on page 10](#)

Prerequisites

- Install sshpass before executing the script file.
 - On Debian/Ubuntu and its derivatives:
 - `sudo apt-get install sshpass`
 - In RedHat/CentOS based systems:
 - `yum install sshpass`
 - `dnf install sshpass` [On Fedora 22+ versions]
- RAM requirement: 4 GiB
- CPU requirement: 4 CPUs
- Storage Class free space requirement: Minimum of 40 GB. I

Configuring HX Profiler for Kubernetes Environments

1. Download the `Cisco-Profiler-2.1-CSI.tar` file from <https://hyperflexsizer.cloudapps.cisco.com>.
 - **DOWNLOAD HX TOOLS > HyperFlex Profiler > Kubernetes**



2. Untar `Cisco-HxProfiler-CSI-4.1.tar.gz` with `tar -zxvf Cisco-HxProfiler-CSI-4.1.tar.gz`.

```
Cisco-HxBench-2.1-CSI      Cisco-HxProfiler-CSI-4.1-original  Cisco-HxProfiler-CSI_old      temp
Cisco-HxProfiler-CSI      Cisco-HxProfiler-CSI-4.1.tar.gz    Cisco-HxProfiler-CSI-old.tar.gz
Cisco-HxProfiler-CSI-4.1  Cisco-HxProfiler-CSI-Final         Cisco-HxProfiler-CSI.tar.gz
cisco@jvdmade-M:~/prabhu$ cd Cisco-HxProfiler-CSI-4.1
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1$ ls
Deployment_yaml_files  Docker_Image  Script_files
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1$
```

3. Run the `cd Cisco-HxProfiler-CSI-4.1/Script_files` to access the directory `Script_files`

```
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1$ ls
Deployment_yaml_files  Docker_Image  Script_files
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1$ cd Script_files/
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$ ls -lrt
total 12
-rw-rw-r-- 1 cisco cisco 10042 Nov 28 13:42 profiler_deployment.sh
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$
```

4. Run the `bash profiler_deployment.sh` script to execute the script.

```
-rw-rw-r-- 1 cisco cisco 10042 Nov 28 13:42 profiler_deployment.sh
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$ ls
profiler_deployment.sh
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$ bash profiler_deployment.sh
```

5. When prompted, type your storage class name and press the enter button.

```
-rw-rw-r-- 1 cisco cisco 10042 Nov 28 13:42 profiler_deployment.sh
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$ ls
profiler_deployment.sh
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$ bash profiler_deployment.sh
-----
Enter storage class name: csi-hxcsi-default
```




Note If you enter wrong storage class name, you will see an error response. If you enter wrong storage name 3 times, the deployment process ends and returns you to the config prompt. Return to step 1 and restart the process.

6. Deploy profiler image in worker nodes: Answer the "Do you want to deploy profiler image in all worker nodes (Recommended) Y/N:" question.
 - If you want to deploy the profiler to all worker nodes, type **Y** and press the **Enter** button and return to Step 8a.
 - If you do not want to deploy image in all worker nodes, type **N** and press the **Enter** button and skip to Step 9.

```
cisco@jvdmade-M:~/profiler_upgraded_rel$ cd Cisco-HxProfiler-CSI-4.1
cisco@jvdmade-M:~/profiler_upgraded_rel/Cisco-HxProfiler-CSI-4.1$ ls
Deployment_yaml_files Docker_Image Script_files
cisco@jvdmade-M:~/profiler_upgraded_rel/Cisco-HxProfiler-CSI-4.1$ cd Script_files/
cisco@jvdmade-M:~/profiler_upgraded_rel/Cisco-HxProfiler-CSI-4.1/Script_files$ ls
profiler_deployment.sh
cisco@jvdmade-M:~/profiler_upgraded_rel/Cisco-HxProfiler-CSI-4.1/Script_files$ bash profiler_deployment.sh
-----
Enter storage class name: csi-hxcsi-default
Do you want to deploy profiler image in all worker nodes(recommended) -Y/N-: y
```

- a. When prompted, type the worker node **username** and password and press the **Enter** button. Make sure that you receive the `Image deployed successfully` message.

```
profiler_deployment.sh
cisco@jvdmade-M:~/prabhu/Cisco-HxProfiler-CSI-4.1/Script_files$ bash profiler_deployment.sh
-----
Enter storage class name: csi-hxcsi-default
Do you want to deploy profiler image in all worker nodes(recommended) -Y/N-: Y
Enter 10.2.14.251 username : cisco
Enter 10.2.14.251 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.251 closed.
Image deployed Successfully
Enter 10.2.14.252 username : █
```

- b. Repeat step 8a to deploy the image in all worker nodes.

```

Enter storage class name: csi-hxcsi-default
Do you want to deploy profiler image in all worker nodes(recommended) -Y/N-: y
Enter 10.2.14.251 username : cisco
Enter 10.2.14.251 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.251 closed.
Image deployed Successfully
Enter 10.2.14.252 username : cisco
Enter 10.2.14.252 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.252 closed.
Image deployed Successfully
Enter 10.2.14.253 username : cisco
Enter 10.2.14.253 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.253 closed.
Image deployed Successfully
The namespace used for deployment is profiler do you want to change -Y/N-: y

```

7. If you do not want to deploy image in all worker nodes, then type **N** and press the **Enter** button. Type IP, username and password of the worker node if you want to deploy image in any specific node.

```

-----
Enter storage class name: csi-hxcsi-default
Do you want to deploy profiler image in all worker nodes(recommended) -Y/N-: n
Enter workernode ip : 10.2.14.252
Enter 10.2.14.252 username : cisco
Enter 10.2.14.252 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.252 closed.
Image deployed Successfully
Do you want to deploy image in any worker node -Y/N-: n

```

8. **Deployment Namespace:** The default namespace used for deployment is `profiler`.
In the provided namespace, it will check if Elasticsearch and postgres pods deployed in the namespace, if yes:

- Delete the entire namespace and deploy fresh namespace, continue to step 11.
- a. If you want to delete and create a new then type **Y** and press the enter button. The default profiler namespace is deleted, and the new namespace named `profiler` is created, type **Y** and press the enter button to confirm the action.

```

-----
Enter storage class name: csi-hxcsi-default
Do you want to deploy profiler image in all worker nodes(recommended) -Y/N-: n
Enter workernode ip : 10.2.14.251
Enter 10.2.14.251 username : cisco
Enter 10.2.14.251 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.251 closed.
Image deployed Successfully
Do you want to deploy image in any worker node -Y/N-: n
The namespace used for deployment is profiler do you want to change -Y/N-: n
Pods Elasticsearch and Postgres are running do you want to delete ns profiler-deploy1 and create them -Y/N-: y
The namespace profiler-deploy1 will be deleted and created freshly do you want to continue -Y/N-: y

```

- b. If you want to deploy in any other namespace type **Y** and type the desired namespace.

```

-----
Enter storage class name: csi-hxcsi-default
Do you want to deploy profiler image in all worker nodes(recommended) -Y/N-: n
Enter workernode ip : 10.2.14.251
Enter 10.2.14.251 username : cisco
Enter 10.2.14.251 password :
Copying docker Images
Loading docker Images
Connection to 10.2.14.251 closed.
Image deployed Successfully
Do you want to deploy image in any worker node -Y/N-: n
The namespace used for deployment is profiler do you want to change -Y/N-: y
namespace must consist of lower case alphanumeric characters or '-', and must start and end with an alphanumeric character
Enter the namespace: profiler-deploy1
Pods Elasticsearch and Postgres are running do you want to delete ns profiler-deploy1 and create them -Y/N-: y
The namespace profiler-deploy1 will be deleted and created freshly do you want to continue -Y/N-: y

```

9. Enter the new namespace name and press the **Enter** button.



- Note** Namespace name requirements: lower case alphanumeric characters and "-" (dash). The name must start and end with an alphanumeric character.

Deployment of Elasticsearch and Postgres yaml file in all nodes are created.

```

The namespace used for deployment is profiler do you want to change -Y/N-: y
namespace must consist of lower case alphanumeric characters or '-', and must start and end with an alphanumeric character
Enter the namespace: new-profiler-test
namespace/new-profiler-test created
configmap/elasticsearch created
persistentvolumeclaim/elasticsearch-pvc created
deployment.apps/elasticsearch created
service/elasticsearch created
persistentvolumeclaim/postgres-pvc created
configmap/postgres-config created
deployment.apps/postgres created
service/postgres created
serviceaccount/mymonitoring created
Error from server (AlreadyExists): error when creating "../Deployment_yaml_files/create_cluster_role.yaml": clusterroles.rbac.authorization.k8s.io "mymonitoring-clusterrole" already exists
clusterrolebinding.rbac.authorization.k8s.io/mymonitoring-clusterrole-binding-new-profiler-test created

```

10. At the prompt for the current UNIX password, type the password for the appadmin GUI account (default password is **password**) and press the **Enter** button.

```

clusterrolebinding.rbac.authorization.k8s.io/mymonitoring-clusterrole-binding-new-profiler-test created
(current) UNIX password:
Enter new UNIX password:
Retype new UNIX password:
Success

```

11. At the prompt for the new UNIX password, type your new user defined password and press the **Enter** button.

```

clusterrolebinding.rbac.authorization.k8s.io/mymonitoring-clusterrole-binding-new-profiler-test created
(current) UNIX password:
Enter new UNIX password:
Retype new UNIX password:
Success

```

12. At the prompt to retype your new UNIX password, type your new password and press the Enter button. Now deployment of profiler will be created.

```
Success
secret/secret-monitoring-auth created
persistentvolumeclaim/profiler-pvc created
deployment.apps/profiler created
service/profiler created
```

- The result is the final profiler-ip with the port number. Copy the IP address and paste it in the browser to fetch the UI

```
-----profiler-ip-----
10.2.14.251:32042
```

Using the Profiler Service

The HX Workload Profiler start and stop services use the `profiler_service.sh` command.


The following table shows the high-level steps for using the profile service.

Task	See
Starting the Profiler Service	Starting the Profiler Service, on page 9
Stopping the Profiler Service	Stopping the Profiler Service, on page 10

Logging in to the Profiler

The HX Workload Profiler user interface (UI) uses system credentials for authentication.

- Step 1** To access the UI, launch a browser window and enter `http://< IP:PORT_NUMBER >` or `http://< IP:PORT_NUMBER >/profiler/index.html` or `http://< IP:PORT_NUMBER >/profiler/index.html`, where the `IP:PORT_NUMBER` is the output of deployment script. The HX Profiler UI appears:

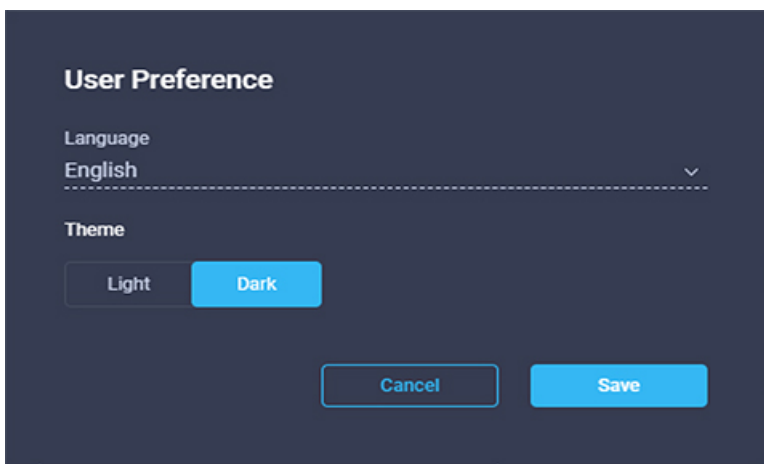


Step 2 When prompted, log in to the UI with the following credentials:

User name: **monitoring**

Password: **<new password set during the deployment workflow>**

Step 3 You can use the User Preference option in the top right corner of the UI to configure **Language** or **Theme**.



Click **Cancel** or **Save** to continue.

Step 4 When finished, you can end the user session by clicking **Logout** at the top right of the page.

Starting the Profiler Service

To start the profiler service run the below commands inside profiler pod:

Run the following command: `sh podscript.sh`.

Stopping the Profiler Service

Complete the following steps inside profiler pod to stop the profiler service:

-
- Step 1** Run the following command: `cd profiler_scripts/`.
- Step 2** Run the following command: `sudo sh profiler_service.sh stop`.
-

Locating the Kubernetes Application Logs

You can find HX Workload Profiler logs in the following locations:

Table 1: Application Logs

Log	Path
Server	/home/monitoring/monitor/server.log
Controller	/home/monitoring/controller/logs/*
Monitor	/home/monitoring/monitor/monitor/monitor.log



CHAPTER 3

Configuring and Using the Profiler Application

- [Adding a Server to the Profiler, on page 11](#)
- [Starting Data Profiling, on page 13](#)
- [Downloading Profiling Results, on page 15](#)
- [Viewing Data Collections from Servers, on page 16](#)

Adding a Server to the Profiler

Adding a Server to the Profiler

At the first login following installation of HX Workload Profiler, you are redirected to the landing page where you can find the + **Add Workload** option on the top right corner of the page. You can select Kubernetes, then you can add multiple Kubernetes nodes.

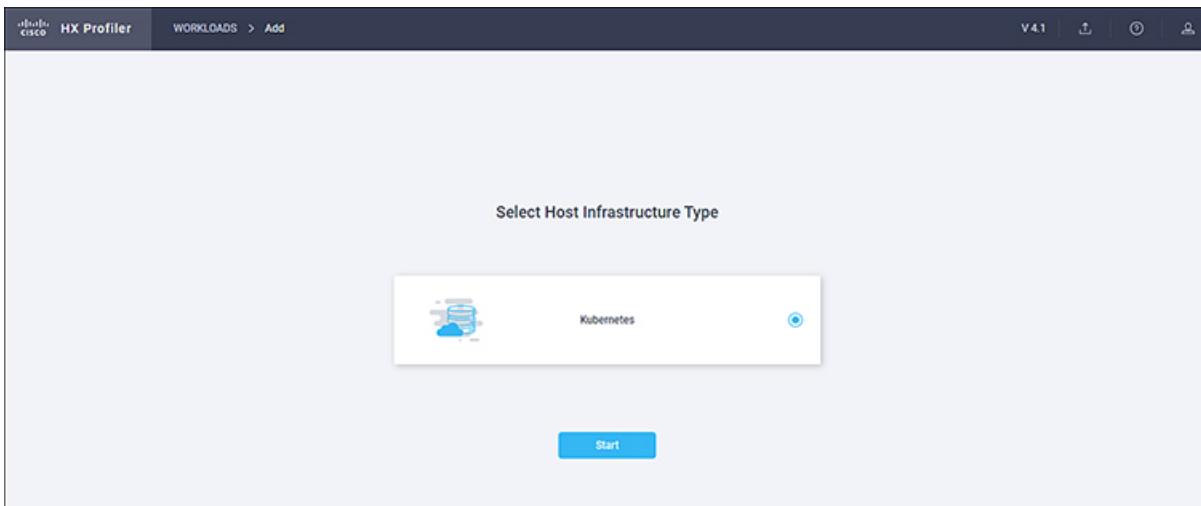
To calculate the metrics for a host, the Profiler captures the metrics for all the Containers on the Nodes. You then need to select the worker nodes to be profiled. By default, all the worker nodes are selected.

The workflow includes:

- **General tab:** Records your Kubernetes node details and then connects to the Kubernetes. If the connection succeeds, the Next button displays enabling you to select the profiling settings.
- **Profiling Setting tab:**, which provides:
 - Profile Name: for specifying the name of the profile which must be unique.
 - Profile period: for specifying the period of profiling. Default value is 7days and the minimum value is 1day.
 - Polling Interval: for specifies the frequency for which polling should occur.
 - Quick Profile: for a quick estimate
 - Detailed Profile: for detailed end-to-end profile results.

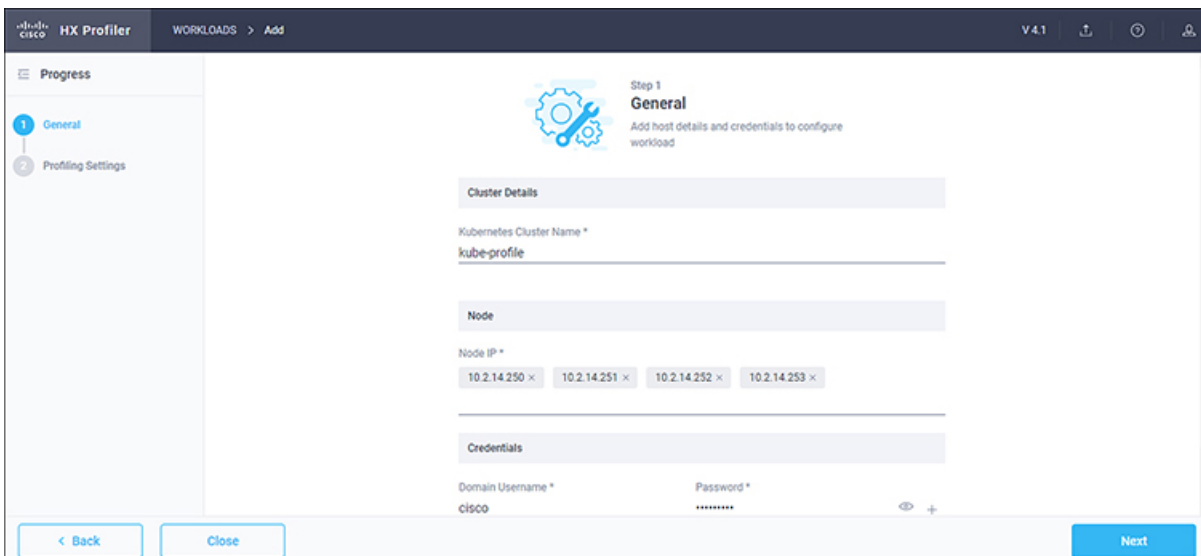
Step 1 Log in to the Profiler. See [Logging in to the Profiler, on page 8](#).

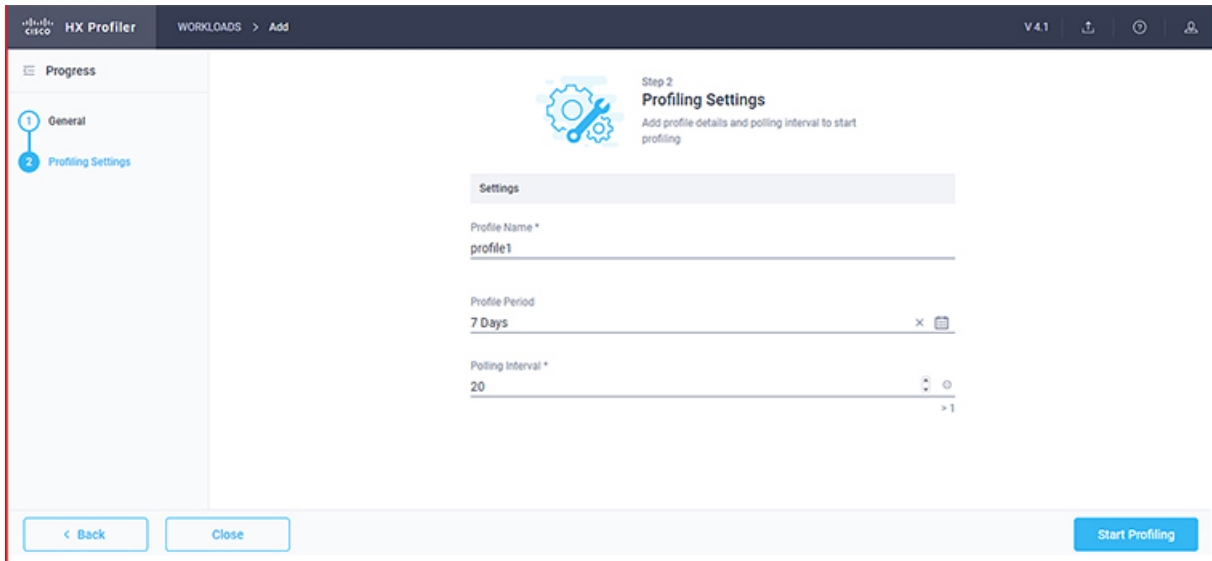
Adding a Server to the Profiler



Step 2 Follow the steps provided to enter values for the following options:

Item	Description
Kubernetes Server Name	Name of the Kubernetes Server you are adding.
User Name	Name of the user as part of the login credentials for the Kubernetes Server.
Password	The password set as part of the login credentials for the Kubernetes Server.
Polling Interval	The interval at which you want polling to occur. The default is 20 seconds. You can change the interval to between 20-120 seconds, based on the number of hosts and Containers being polled in that Kubernetes Server.





- Step 3** Click **Start Profiling**. The Profiler service starts automatically. When the connection status indicates successful connection, click **Next**.
- Step 4** On the **Poll filter** tab, review the hierarchy values.

Starting Data Profiling

Starting Kubernetes Data Profiling

Following the successful addition of a Kubernetes Server, the new Kubernetes Server displays on the Data Inventory (home) page. You configure the profiling attributes by providing the profile name and duration.

Table 2: Profiling Operation Options

Item	Description
Delete (trash can)	Deletes a previously added Kubernetes Server.
Edit (pen)	Edit Kubernetes Server properties to add or remove hosts for polling.
Stop (symbol)	Stops the profiling so you can resume it later.
Reset (refresh symbol)	Performs a reset operation, which creates a new profile and starts polling. When you trigger reset, the profiler stops the active/running profile and creates a new one. A prompt asks for confirmation.
View Collection	Opens the View Collection page so you can browse through the collected data as part of the profiling to review the HOST and Container level data.

Polling starts as soon as you enter the profiler name, days, and polling period.

Step 1 In the dialog that displays, perform the following steps:

- a) Enter a name for the profile.
- b) Select a duration value from the **Profiling Period** down-drop list.
- c) Click **Ok**.

The screenshot shows the 'Profiling Settings' dialog in the Cisco HX Profiler application. The dialog is titled 'Step 2 Profiling Settings' and includes the instruction 'Add profile details and polling interval to start profiling'. The 'Settings' section contains three input fields: 'Profile Name *' with the value 'profile1', 'Profile Period' with the value '7 Days', and 'Polling Interval *' with the value '20'. At the bottom of the dialog, there are three buttons: '< Back', 'Close', and 'Start Profiling'.

The screenshot shows the 'Kubernetes kube-profile' dashboard in the Cisco HX Profiler application. The dashboard displays the following information:

- Node Reachability:** 4 Reachable (green), 0 Not Reachable (red).
- Profiling Summary:** 4882 SUCCESSFUL (green bar).
- Profiling Status:** In Progress (47% complete), Remaining Duration: 3d 17h 18m.
- Profile Details:**
 - Profile Name: kube-profile
 - Interval: 20 Sec
 - Total Duration: 7d
 - Elapsed Duration: 3d 6h 41m
 - Initial Start Time: Dec 17, 2021 04:41 PM

Buttons for 'View Collection' and 'Download' are visible on the right side of the dashboard.

Following successful profile creation, the Profiler begins polling the selected hosts and all the Containers on those hosts. When the polling starts, the data collector runs as a background process. The Datacenter Inventory page displays information about the hosts and polling, showing number of hosts and the status of the polling.

Step 2 (Optional) To stop an in-progress profiling operation so you can resume it later, see [Stopping the Profiler Service, on page 10](#).

Step 3 (Optional) To stop an in-progress profiling operation, click **Reset**. You can then start a new one.

Step 4 (Optional) To browse through the collected data, see [Viewing a Kubernetes Collection, on page 20](#).

Step 5 (Optional) To download profiling data, see [Downloading Kubernetes Profiling Results, on page 15](#).

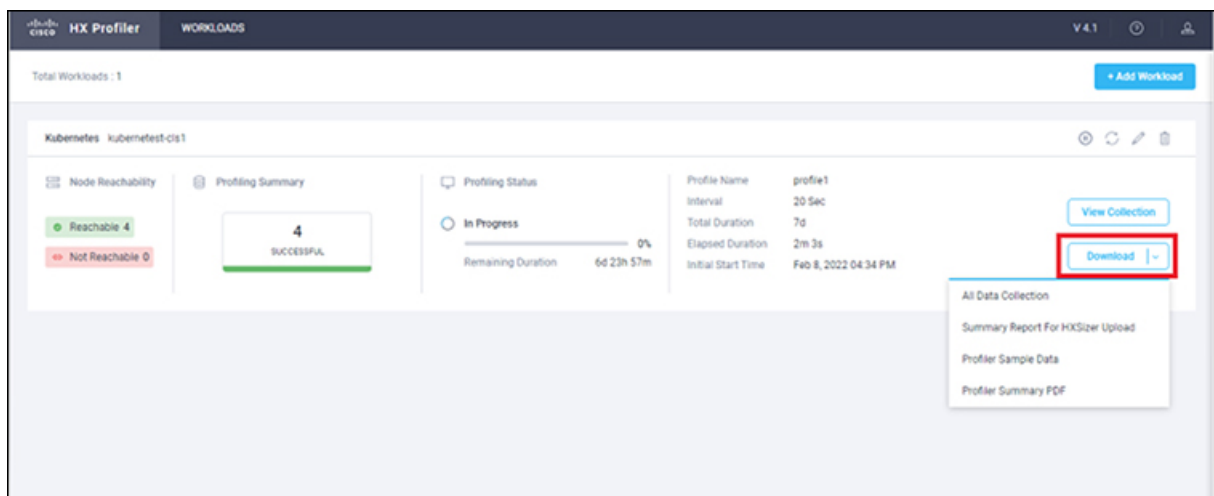
Downloading Profiling Results

Downloading Kubernetes Profiling Results

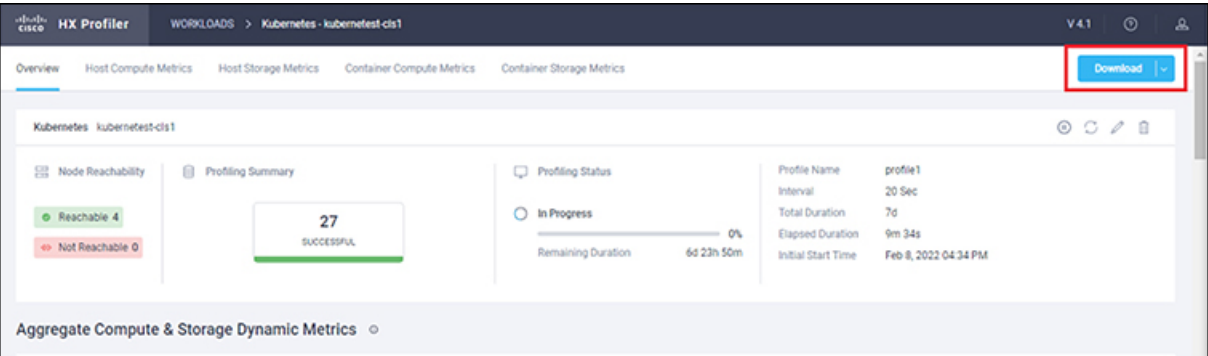
Step 1 On the Datacenter Inventory page, locate and select the profile whose data you want to download.

Step 2 Use the **Download** option to select one of the following:

Option	Description
All Data Collection	Downloads summary, time series, and CVS for both the host and Container
Summary Report for HxSizer Upload	Collects information from Kubernetes cluster. The download provides the output in CSV format and can be directly uploaded to the Compute and Capacity Workload of HxSizer.
Profiler Sample Data	Downloads the sampled data for the selected profile in the following formats: <ul style="list-style-type: none"> Summarized host data (CVS) Summarized Container data (CVS) Time series data of host (zipped CSV file) Time series data of Container (zipped CSV file)
Profiler Summary PDF	PDF download



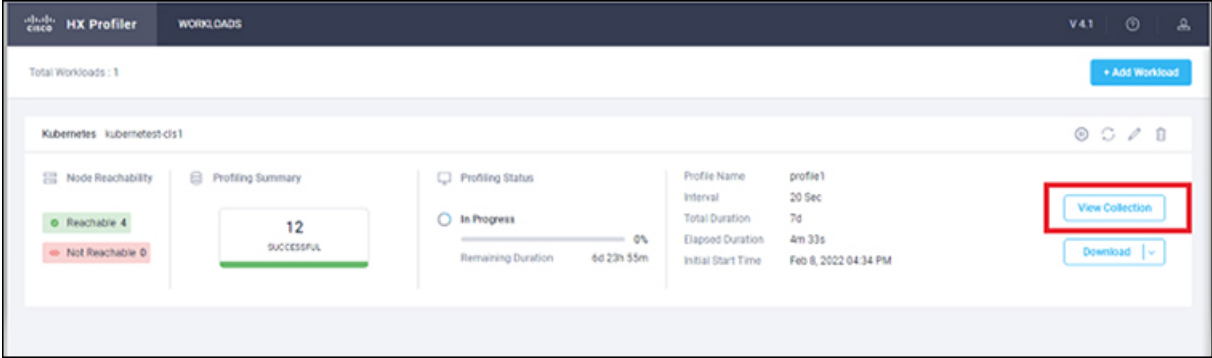
Step 3 You can still download the profiling results when viewing the compute, storage, and network data of various hosts and Containers by clicking on the **Download** button on the top right corner of the UI.



Viewing Data Collections from Servers

Viewing Data Collections from Kubernetes Server

The View Collection page has five tabs at the top left of the page, Overview, Host Compute Metrics, Host Storage Metrics, Container Compute Metrics, Container Storage Metrics. When clicked, they show summary data described in this section. The View Collection page also provides the summarized data for Compute Summary and Storage Summary of individual node and Containers. You can also fetch the data for specific period of profiling using the predefined filter present on the top right corner of the page with the minimum being 30minutes or the user can also use the timeline to select the time period. The following sections describe summary data that appears in each tab and view through filter tool use.



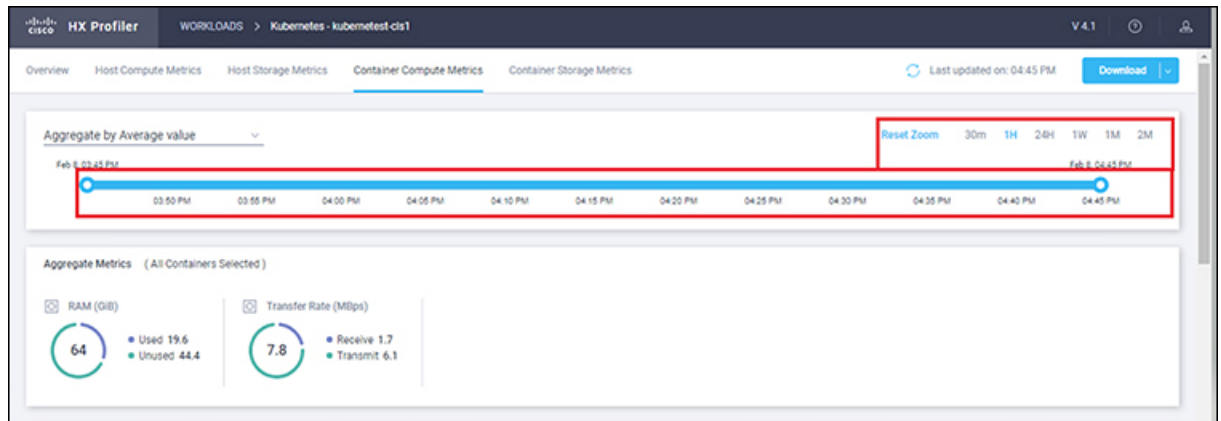
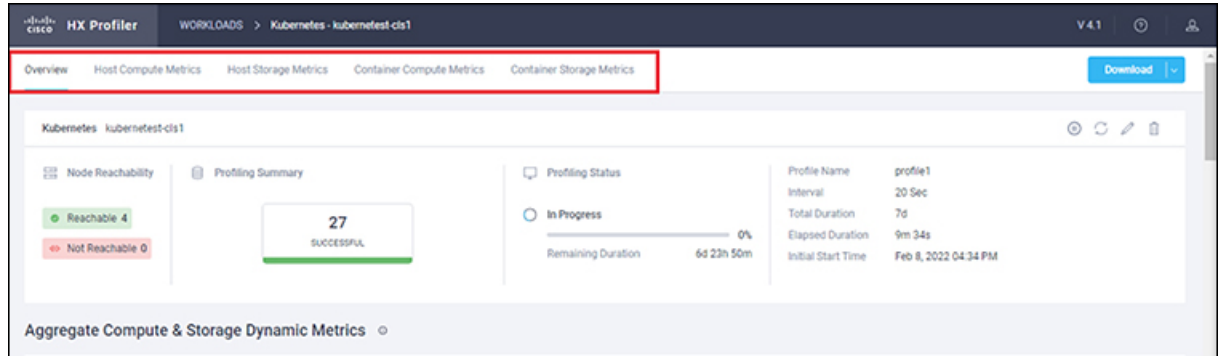
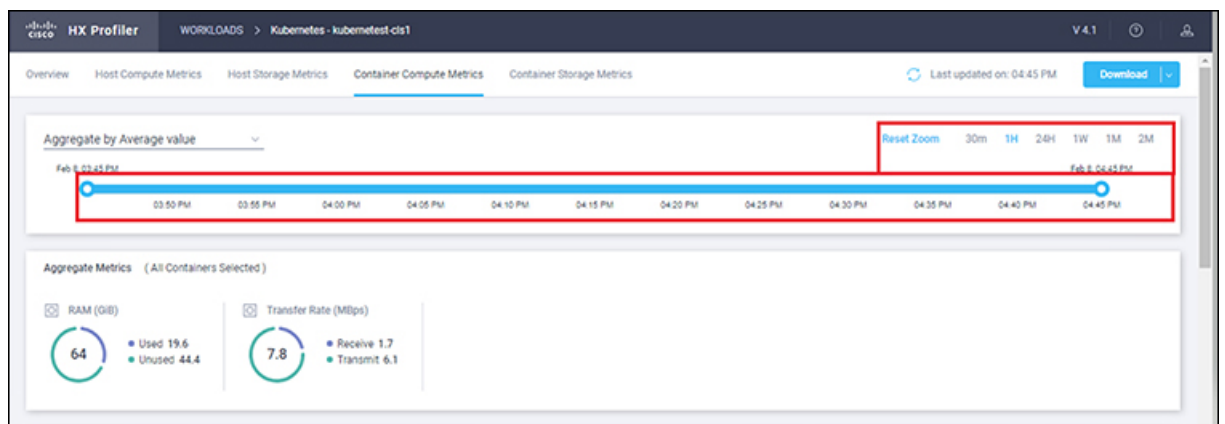
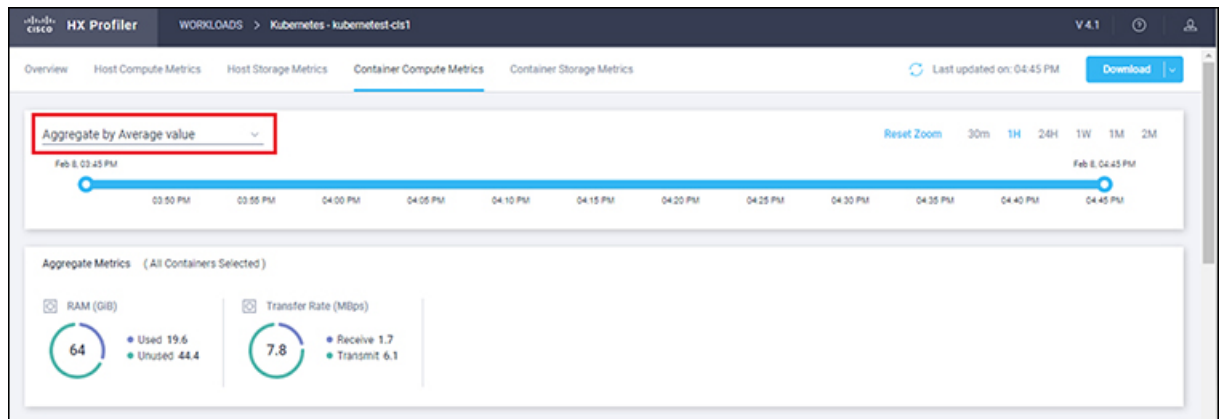


Table 3: Host View Filter Options

Item	Description
Aggregation	Filters to view the summarizations based on peak or average selections. Your selection determines the display of the table metrics and trends. The summarized values represent the following: <ul style="list-style-type: none"> • Peak: Peak value of all the metrics in the selected interval • Average: Averages of all the values in the selected interval

You can filter Host metrics using the search option by host name only. For Container metrics, you can filter with the Container name, using the search option.



Host and Container Summarized Metrics

Based on your time period and parameters selections, the Profiler computes and populates the compute and storage metrics. Metrics display for the following values:

Host View Compute Table

- Host_Name
- Processor Type
- Clock (GHz)
- #Physical Cores
- CPU Util (%)
- RAM (GB)
- RAM Util (GB)
- Network Throughput- Rx (MBps)
- Network Throughput- Tx (MBps)

Host View Storage Table

- Provisioned Capacity (TB)

- Used Storage Capacity (TB)
- Read Throughput (MBps)
- Write Throughput (MBps)
- Read (%)
- Write (%)
- Read IOPS
- Write IOPS
- Read Latency (ms)
- Write Latency (ms)

Container View Compute Table

- Container Name
- Clock (GHz)
- #Physical Cores
- CPU Util (%)
- RAM (GB)
- RAM Util (GB)
- Network Throughput- Rx (MBps)
- Network Throughput- Tx (MBps)

Container View Storage Table

- Container Name
- Disk Capacity (GB)
- Disk Used (GB)
- Read Throughput (MBps)
- Write Throughput (MBps)
- Read (%)
- Write (%)
- Read IOPS
- Write IOPS
- Read Latency (ms)
- Write Latency (ms)

Host and Container Trends

The View Collection page provides trend charts and an overview of various parameters at a host level and Container level for both compute and storage parameters. The overview provides information aggregate storage and compute matrix along with a 30-day sizing summary. You can view the trend charts by selecting the host or the Container from the table.

Metrics for the following values display in the trend charts:

Host View Compute Trends

- RAM Utilization (GiB)
- Receive Rate (MBps)
- Transmit Rate (MBps)

Container View Compute Trends

- RAM Utilization (GB)
- Receive Rate (Mbps)
- Transmit Rate (Mbps)



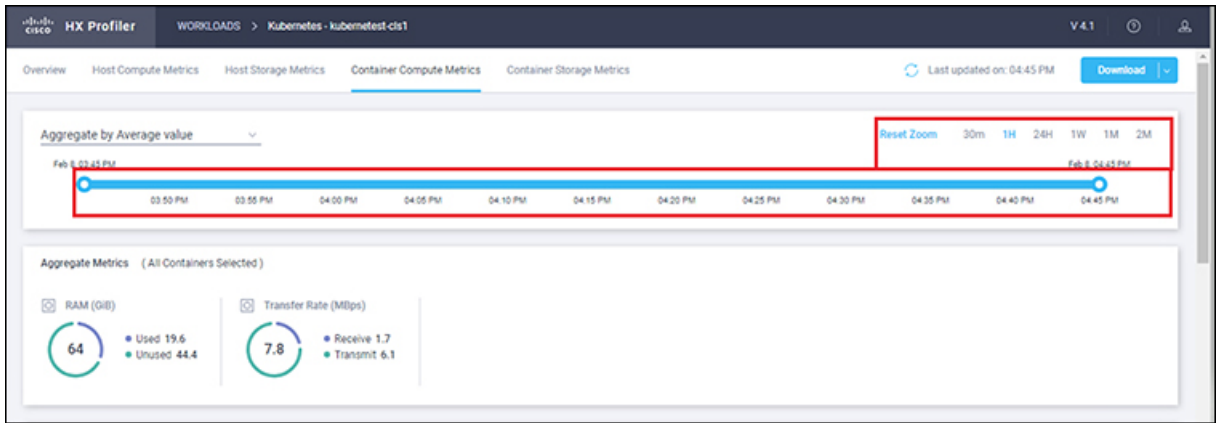
Host and Container View Storage Trends

- Read Throughput (MBps)
- Write Throughput (MBps)
- Read Ratio
- Write Ratio
- Read IOPS
- Write IOPS
- Read Latency (ms)
- Write Latency (ms)

Viewing a Kubernetes Collection

You can also use the filter and search tool in the Hosts and Container views to display only those Hosts and Containers that you want to see.

-
- Step 1** In the Datacenter Inventory page, click **View Collection** to browse through the collected data.
- Step 2** In the View Collection page, select between **Host Compute Metrics**, **Host StorageMetrics**, **Container Compute Metrics** or **Container Storage Metrics**.
- Step 3** In the **Select Containers** tab, select or unselect the toggle switch for the Containers you want to view, then click **Next**. All Containers are selected by default.
- Step 4** To view data for a different time period, select from the options in the top right corner with the default minimum value of 30 minutes. Use the time slider above the fixed time period selection to view data for a specific time range from the selected time period.





CHAPTER 4

Monitoring Approach

- [Download Quick Profile \(30-Days\) Kubernetes Host Summary Statistics and Data](#) , on page 23
- [Profiler Data Collection Architecture](#), on page 23
- [Performing Profiler Clean-up](#), on page 24

Download Quick Profile (30-Days) Kubernetes Host Summary Statistics and Data

You can download the Quick Profile (30-days) Kubernetes Host Summary statistics/data report from the Profiler Home page > **Download Results** option.

Report characteristics:

- Includes Container/host level compute and storage capacity metrics
- Does not include deep storage or compute metrics or metrics for individual Containers
- Data downloads in CSV format and can be uploaded to the Compute and Capacity workload in the HX Sizer tool

Profiler Data Collection Architecture

Understanding the Kubernetes Profiler Data Collection Architecture

The Profiler connects to Kubernetes server using the root admin credentials.

Architecture characteristics:

- The Profiler directly collects the Kubernetes server Inventory information (Clusters, Hosts, Containers, and basic metrics) and the compute metrics from Kubernetes server itself.
- The Profiler also collects deep storage metrics for each Container and aggregates these metrics at a host-level.
- The Kubernetes server does not directly maintain deep storage metrics. The Profiler software invokes Kubernetes server to fetch this data from each individual host (using vSCSI) to obtain detailed data for

each Container. The Profiler then summarizes the data from all Containers to build a host-level summary. For example, the software obtains maximum, minimum, and average values for metrics, such as latency, where as IOPS (Total Blocks/Interval) and throughput (Total Bytes/Interval) are average values only.

Performing Profiler Clean-up

Performing Kubernetes Profiler Container Clean-up

After you complete your profiling activities, follow the best practice of performing the Profiler Pod clean-up to achieve a thorough shutdown and exit of the Profiler.

-
- Step 1** Download the profiled data with the following steps:
- a) Launch the Profiler application.
 - b) Download the data (CSV format) from the Home page. See [Downloading Kubernetes Profiling Results, on page 15](#).
 - c) Save the CSV for further analysis.
- Step 2** Delete the namespace if it is no longer required to profile the data for the environment.
-



CHAPTER 5

Troubleshooting

- [Troubleshooting, on page 25](#)

Troubleshooting

Troubleshooting Section for Kubernetes

1. If the selected hosts are not reachable, then perform the following steps:
 - Check the connection status of the hosts in Kubernetes.
 - Make sure all the selected hosts connection status is **Ready**.
2. If you forgot the Profiler Pod password, then perform the following step:
 - Delete the old deployed namespace and redeploy the new one.

