

IBM Hyper-Scale Manager
Version 5.5.3

User Guide



Note

Before using this information and the product it supports, read the information in [Notices](#).

First Edition (June 2021)

This edition applies to IBM Hyper-Scale Manager version 5.5.3 and to all subsequent releases and modifications, until otherwise indicated in new editions. Newer document editions may be issued for the same product version in order to add missing information, update information, or amend typographical errors.

© **Copyright International Business Machines Corporation 2013, 2021.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

| | |
|--|------------|
| Figures..... | V |
| About this guide..... | vii |
| Who should use this guide..... | vii |
| Conventions used in this guide..... | vii |
| Related information and publications..... | vii |
| Getting information, help, and service..... | viii |
| IBM Publications Center..... | viii |
| Feedback..... | viii |
| Chapter 1. Introducing IBM Hyper-Scale Manager..... | 1 |
| Chapter 2. Installation..... | 3 |
| Installing IBM Hyper-Scale Manager..... | 3 |
| Completing initial setup..... | 5 |
| Upgrading IBM Hyper-Scale Manager version 5.x..... | 5 |
| Uninstalling IBM Hyper-Scale Manager version 5.x..... | 7 |
| Chapter 3. Configuration and administration..... | 9 |
| Migrating IBM Hyper-Scale Manager inventory from XIV Management Tools 4.x | 9 |
| Changing the display language..... | 9 |
| Adding a system after initial setup..... | 12 |
| Configuring an access code..... | 15 |
| Replacing the IBM Hyper-Scale Manager certificate..... | 16 |
| Trusting the IBM Hyper-Scale Manager certificate..... | 18 |
| Backing up and restoring IBM Hyper-Scale Manager..... | 19 |
| Backing up IBM Hyper-Scale Manager..... | 19 |
| Restoring IBM Hyper-Scale Manager..... | 20 |
| Managing connectivity..... | 21 |
| Defining new connectivity..... | 21 |
| Chapter 4. Viewing the dashboard and system health score..... | 27 |
| Chapter 5. Understanding and monitoring capacity..... | 31 |
| Monitoring your capacity..... | 33 |
| Chapter 6. Planning capacities..... | 37 |
| Generating a capacity planning report..... | 38 |
| Capacity planning report structure..... | 38 |
| Monitoring capacity statistics..... | 40 |
| Understanding deduplication-aware capacity management calculations..... | 41 |
| Deactivating IBM Hyper-Scale Manager deduplication-aware capacity management calculations..... | 43 |
| Moving capacity data between Management Server instances..... | 44 |
| Exporting the raw capacity data..... | 44 |
| Importing capacity data..... | 44 |
| Resetting the raw capacity data..... | 45 |
| Chapter 7. Troubleshooting latency..... | 47 |

| | |
|----------------------|-----------|
| Notices | 49 |
| Trademarks..... | 50 |
| Index | 51 |

Figures

- 1. Changing the display language..... 10
- 2. Language Preferences..... 11
- 3. Restarting the Hyper-Scale Manager..... 12
- 4. Adding a system to the inventory..... 13
- 5. Adding a destination email.....14
- 6. Advanced login screen..... 16
- 7. Viewing system certificate, licensing and authentication..... 18
- 8. Creating new connectivity..... 22
- 9. System Connectivity view..... 23
- 10. Add IP interface form..... 24
- 11. System Connectivity details view..... 25
- 12. Health widget on the Dashboard..... 27
- 13. Configurable tips..... 29
- 14. Projected vs. Allocated vs. Physical31
- 15. Projected System Capacity tool..... 32
- 16. Dashboard..... 33
- 17. Select tier..... 34
- 18. Capacity view..... 35
- 19. System allocation over time..... 39
- 20. System by usage growth rate.....39
- 21. System usage - detailed graphs..... 40
- 22. Illustration of the latency breakdown..... 47
- 23. Latency view from Hyper-Scale Manager..... 48

About this guide

This user guide describes how to install, configure, and use IBM Hyper-Scale Manager for managing multiple supported storage systems.

Who should use this guide

This user guide is intended for storage administrators who intend to manage one or more supported IBM storage systems through the IBM Hyper-Scale Manager server and graphical user interface.

Basic familiarity with the following is required:

- Linux[®] operating system
- Network, server, and storage terms
- IBM storage system to be managed with IBM Hyper-Scale Manager

Conventions used in this guide

These notices are used in this guide to highlight key information.

Note: These notices provide important tips, guidance, or advice.

Important: These notices provide information or advice that might help you avoid inconvenient or difficult situations.



Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage can occur.

Related information and publications

Additional information and publications related to IBM Hyper-Scale Manager can be found on the following information sources.

- [IBM FlashSystem[®] documentation](http://ibm.com/docs/en/hyper-scale-manager) (ibm.com/docs/en/hyper-scale-manager)
- [IBM FlashSystem A9000 documentation](http://ibm.com/docs/en/flashsystem-a9000) (ibm.com/docs/en/flashsystem-a9000)
- [IBM FlashSystem A9000 documentation](http://ibm.com/docs/en/flashsystem-a9000r) (ibm.com/docs/en/flashsystem-a9000r)
- [IBM XIV[®] Storage System documentation](http://ibm.com/docs/en/xiv-storage-system) (ibm.com/docs/en/xiv-storage-system)
- [IBM Spectrum Accelerate documentation](http://ibm.com/docs/en/xiv-spectrum-accelerate) (ibm.com/docs/en/xiv-spectrum-accelerate)
- [IBM IT Infrastructure](http://ibm.com/it-infrastructure) (ibm.com/it-infrastructure)
- [IBM Storage Redbooks[®] website](http://redbooks.ibm.com/portals/storage) (redbooks.ibm.com/portals/storage)

IBM Hyper-Scale Manager documentation set

IBM Hyper-Scale Manager 5.x currently has the following documentation set.

| IBM Hyper-Scale Manager publications | Description |
|--------------------------------------|--|
| <i>Release Notes</i> | Describes what's new, requirements, compatibility, change log, and known issues information in the latest version. |

| IBM Hyper-Scale Manager publications | Description |
|--------------------------------------|--|
| <i>User Guide</i> | Describes how to install, configure, and operate IBM Hyper-Scale Manager. |
| <i>Quick-Start Guide</i> | Walks you through installing the IBM Hyper-Scale Manager software and initial configuration of the GUI. |
| <i>REST API Specifications</i> | Provides specifications of the IBM Hyper-Scale Manager Representational State Transfer (REST) application programming interface (API). |

Getting information, help, and service

If you need help, service, technical assistance, or want more information about IBM products, you can find various sources to assist you. You can view the following websites to get information about IBM products and services and to find the latest technical information and support.

- [IBM website \(ibm.com®\)](http://ibm.com)
- [IBM Support Portal website \(www.ibm.com/storage/support\)](http://www.ibm.com/storage/support)
- [IBM Directory of Worldwide Contacts website \(www.ibm.com/planetwide\)](http://www.ibm.com/planetwide)

IBM Publications Center

The IBM Publications Center is a worldwide central repository for IBM product publications and marketing material.

The [IBM Publications Center website \(www.ibm.com/shop/publications/order/\)](http://www.ibm.com/shop/publications/order/) offers customized search functions to help you find the publications that you need. You can view or download publications at no charge.

Sending comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this publication or any other IBM storage product documentation:

- Send your comments by email to ibmdocs@us.ibm.com. Be sure to include the following information:
 - Exact publication title and version
 - Publication form number (for example, GA32-1234-00)
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Chapter 1. Introduction

IBM Hyper-Scale Manager provides a single and centralized management platform for multiple and different members of the IBM Spectrum Accelerate Family, including **IBM FlashSystem A9000** and **A9000R**, **IBM XIV Gen3** and **IBM Spectrum Accelerate** software-defined storage (SDS) solutions.

IBM Hyper-Scale Manager reduces operational complexity and enhances capacity planning through integrated management for large storage system deployments. Storage administrators can log in to the Management Server and employ the advanced web-based IBM Hyper-Scale Manager to configure, manage, and monitor multiple storage systems.

Easy-to-use, browser-based user interface

IBM Hyper-Scale Manager provides a breakthrough user experience with powerful health monitoring that integrates events and alerts across the managed systems. The IBM Hyper-Scale Manager UI offers context-based interaction on a single page with advanced filtering and navigation capabilities, as well as advance statistics monitoring.

You don't need to be an expert to work with IBM Hyper-Scale Manager. Just log in, run the intuitive "first-time" configuration, and get started working.

Note: For detailed information about browser and screen resolution requirements for using IBM Hyper-Scale Manager on client PCs, see the Hyper-Scale Manager release notes.

Cloud integration component

From version 5.0.x, IBM Hyper-Scale Manager includes a new component, cloud integration, which is based on IBM Spectrum® Connect (previously known as IBM Spectrum Control Base Edition). IBM Spectrum Connect is a centralized server system that consolidates a range of IBM storage provisioning, automation, and monitoring solutions through a unified server platform. It provides a single server back end location and enables centralized management of IBM storage resources for different virtualization, cloud and container platforms.

Implementing sophisticated storage abstraction techniques, IBM Spectrum Connect helps storage and other infrastructure administrators easily provision and integrate IBM storage with VMware solutions, Kubernetes persistent volumes, and Microsoft PowerShell automation methods.

Note: The cloud integration function can be applied only to storage systems that are managed using IBM Hyper-Scale Manager 5.0.x or later. The cloud integration component is accessible from the cloud integration hyperlink on the Hyper-Scale Manager login page. For more information about IBM Spectrum Connect, see [IBM Spectrum Connect on the Knowledge Center](#).

System management through REST API

Programmable automation and control of your storage systems can be implemented through the IBM Hyper-Scale Manager REST API.

Availability features

Running more than one Hyper-Scale Manager allows you to ensure cross-system availability. There are two ways to achieve availability for Hyper-Scale Manager:

- Manually, by deploying 2 instances (primary and standby) of Hyper-Scale Manager.

In case the primary Hyper-Scale Manager host fails, the primary Hyper-Scale Manager setup and configuration can be restored on the standby Hyper-Scale Manager. (See [“Backing up and restoring IBM Hyper-Scale Manager”](#) on page 19 for more information.)

The browser can then be pointed to the standby Hyper-Scale Manager which is now active.

- Using virtualization technologies.

Virtualization technologies (such as VMware tools) can be used to achieve high availability.

Chapter 2. Installation

Installation and initial setup of IBM Hyper-Scale Manager using the web-based user interface (UI) is fast and easy. If you are a Storage Administrator, these instructions are intended for you.

Installing IBM Hyper-Scale Manager

Before you begin working with IBM Hyper-Scale Manager, install the Management Server software.

Before you begin

Ensure that your system meets all system requirements. For detailed information, refer to the Hyper-Scale Manager release notes.

Note: If not all system requirements can be met for the deduplication-aware capacity management calculations, they can be deactivated post-installation. For more information, see [“Deactivating IBM Hyper-Scale Manager deduplication-aware capacity management calculations”](#) on page 43.

About this task

Follow the steps to install Hyper-Scale Manager.

Procedure

1. Download the installation file from Fix Central. The name of the installation file should be similar to: IBM-Hyper-Scale-Manager_<version>.bin.
2. FTP the downloaded installation package file to your Hyper-Scale Manager server by using the root username.
3. Log in to your Hyper-Scale Manager server and change the directory to the one that contains the installation package file.
4. Make sure that run permissions are set.
5. Run the following command:

```
$ chmod +x IBM-Hyper-Scale-Manager_<version>.bin
```

6. Run the following command:

```
./IBM-Hyper-Scale-Manager_<version>.bin
```

A message similar to the following is displayed.

```
You are about to install IBM Hyper-Scale Manager in the [/home/msms] directory.  
Permissions to this directory are designated only to the 'msms' user created during  
installation.  
Keep in mind that both the installation and '/opt' directories require directory  
traversal permissions (x-permission) for the 'msms' user.  
Press Enter to accept the installation location, or type a different path to change the  
installation directory. >
```

7. Press Enter to accept the installation location or type a different path to change the location and then press Enter.

If you decide to type a different path to change the installation location, and there are no traversal permissions (x-permission) to the directory, a message similar to the following is displayed.

```
Directory [installation location] permissions currently do not allow traversal
for 'msms' user. Do you want to add traversal permissions (x-permission) to
allow access? [Y/N] >
```

8. Type `y` to allow access.

Traversal permissions (x-permission) are added to the selected directory.

Note: If you decide not to add traversal permissions (x-permission), you must type a different path in order to continue with the installation.

Note: If you don't have traversal permissions (x-permission) to the `/opt/` directory for the cloud integration feature, a message similar to the one shown in step 6 is also displayed.

A message similar to the following is displayed.

```
Enter HTTP port number to use with IBM Hyper-Scale Manager or 0 to disable
HTTP port [8080] >
Enter HTTPS port number to use with IBM Hyper-Scale Manager [8443] >
Enter HTTPS port number to use with cloud integration feature [8440] >
```

9. Enter the communication ports at the prompts and press `Enter`.

A message similar to the following is displayed.

```
IBM-Hyper-Scale Manager version <version> will be installed under /home/msms/hyperscale.
Are you sure you want to install now? [y/n] >
```

10. Enter `y` to begin the installation process.

The license agreement is displayed.

11. Press `1` and then `Enter` to accept the license and continue with the installation.

Note: In our constant efforts to improve your experience with Storage Management, we added an analytics module to help us understand how you and others are using the application.

During the installation, you are asked to approve activating this analytics module. No private information is shared with IBM, including names, sizes, or passwords. If you agree, only information about the way you interact with the application is sent to IBM and treated as confidential.

What to do next

When the installation is complete, make sure that all the required ports are opened on the OS firewall:

- Outbound:
 - TCP 25 (SMTP)
 - TCP 7778 (back-end communication with storage systems)
 - Inbound (default):
 - Web UI port - 8080 (for HTTP support) or 8443 (for HTTPS support)
 - Cloud integration port - 8440
-

Note: For IPv6 connectivity, make sure the IPv6 firewall (the `ip6tables` in RHEL 6) is configured to allow traffic to these ports.

Reboot your system and make sure that the ports are correctly configured on your firewall.

Go to [“Completing initial setup”](#) on page 5.

Completing initial setup

Once you have installed IBM Hyper-Scale Manager, run the first-time login and setup wizard to configure your environment.

Before you begin

Connect and log in to the IBM Hyper-Scale Manager UI on the web with one of the following URLs:

- `http://<hostname>:8080` or `http://<IP address>:8080` (HTTP supported browser access)
- `https://<hostname>:8443` or `https://<IP address>:8443` (HTTPS supported browser access)

The first-time login and setup wizard is displayed in your browser.

Procedure

1. From the **Welcome to IBM Hyper-Scale Manager** page, click **Let's Start**.

Note: You'll need the IP/DNS of one storage system (for example, IBM XIV Gen3, IBM Spectrum Accelerate, IBM FlashSystem A9000 or IBM FlashSystem A9000R system) and the Storage Administrator credentials to authenticate to that system.

2. Add a storage system to your inventory. Enter the IP/DNS address of the system and click **Continue**.
3. Log in to the system with your Storage Administrator credentials and click **Continue**. Use the default username **admin** and the default password **adminadmin**, unless the defaults have been previously changed. In this case, log in with the new username and password.

Note: If you are logging in to a new or upgraded system, a Certificate Error is displayed. First, trust the certificate and then click **Continue**.

4. At this point, you need to configure the Management Server. IBM Hyper-Scale Manager uses the Management Server as a mediator to simultaneously monitor and configure multiple systems in your inventory. Click **Continue**.
5. Create a monitoring account for the Management Server. If multiple Hyper-Scale Manager servers are installed, and possibly manage the same storage systems, make sure to use the same user credentials for the monitoring account on **all** of your Management Servers.
 - If a monitoring account hasn't been generated previously, then a username and password are automatically generated for you.
 - If a monitoring account has been generated previously, you can set a new monitoring account name and password. Keep in mind that predefined usernames, such as *admin*, can't be used for monitoring accounts.
 - If you are using an LDAP domain, enter the name and password that exists in the LDAP server. If your LDAP server requires a domain, use UPN notation (for example, *name@domain.ibm.com*).

Note: If you are logging in to a new or upgraded system, a *Software License Agreement* is displayed. Click **Accept** and **Continue**.

6. **Congratulations! You're Done.** Now that you've completed the installation and setup, click **Show me the Dashboard** and get started managing your storage.

Upgrading IBM Hyper-Scale Manager version 5.x

Upgrade IBM Hyper-Scale Manager version 5.x to the most recent version available.

Before you begin

Ensure that your system meets all system requirements. For detailed information, refer to the Hyper-Scale Manager release notes.

Note: If not all system requirements can be met for the deduplication-aware capacity management calculations, they can be deactivated post-upgrade. For more information, see [“Deactivating IBM Hyper-Scale Manager deduplication-aware capacity management calculations” on page 43.](#)

Important: Before upgrading from version 5.5.1 of IBM Hyper-Scale Manager, be sure to first restart Hyper-Scale Manager and then immediately complete the upgrade process.

Failure to do so may result in an unstable upgrade and server state.

To restart:

1. Go to **Settings > Management Server.**
 2. From the Actions menu select **Restart > Restart Hyper-Scale Manager.**
-

About this task

Follow the steps to upgrade IBM Hyper-Scale Manager and the Management Server from version 5.x to the current version.

Procedure

1. Start by downloading the installation file from Fix Central. The name of the installation file should be similar to: `IBM-Hyper-Scale-Manager_<version>.bin`.
2. FTP the downloaded installation package file into your Management Server server by using the root username.
3. Log in to your Management Server and change the directory to the directory that contains the installation package file.
4. Make sure that run permissions are set. Run the following command:

```
$ chmod +x IBM-Hyper-Scale-Manager_<version>.bin
```

5. Run the following command:

```
./IBM-Hyper-Scale-Manager_<version>.bin
```

A menu similar to the following is displayed:

```
-----  
-----  IBM Hyper-Scale Manager <version>  -----  
-----  
IBM Hyper-Scale Manager version <version> is already installed on this server.  
Please select one of the following:  
1) Upgrade  
2) Uninstall  
3) Exit  
Please enter your choice
```

6. Type **1** to upgrade and press Enter.
A message similar to the following is displayed:

```
Are you sure you want to upgrade now? [y/n]
```

7. Type **y** to begin the upgrade.
A prompt to view the license agreement is displayed.

IBM Hyper-Scale Manager license agreement will be shown.
Press 'Enter' to view the license agreement.
During the review of the license, you can press '1' and 'Enter' to accept and continue with the installation, or '2' and 'Enter' to decline and abort the installation.

8. Press Enter to view the license agreement.
9. Type 1 and then Enter to accept the license and continue with the IBM Hyper-Scale Manager installation.

Note: In our constant efforts to improve your experience with Storage Management, we added an analytics module to help us understand how you and others are using the application. During the installation, you are asked to approve activating this analytics module. No private information is shared with IBM; including names, sizes, or passwords. If you agree, only information about the way you interact with the application is sent to IBM and treated as confidential.

10. If you agree to allow access to your IBM Hyper-Scale Manager usage data, type y. If you don't want to allow access to your IBM Hyper-Scale Manager usage data, type n.

Note: Access to your IBM Hyper-Scale Manager usage data does not include personal data.

The upgrade procedure begins.

11. Continue working when the upgrade procedure completes successfully.

Uninstalling IBM Hyper-Scale Manager version 5.x

Use this information to help you uninstall IBM Hyper-Scale Manager version 5.x.

Before you begin

Important: The installation contains VMware integration data that will be lost during the uninstall. Be sure to back up your cloud integration configuration before uninstalling IBM Hyper-Scale Manager.

About this task

Follow the steps to uninstall IBM Hyper-Scale Manager and the Management Server version 5.x.

Procedure

1. Start by downloading the installation file from Fix Central. The name of the installation file should be similar to: IBM-Hyper-Scale-Manager_<version>.bin.
2. FTP the downloaded installation package file into your Management Server server by using the root username.
3. Log in to your Management Server server and change the directory to the directory that contains the installation package file.
4. Make sure that run permissions are set. Run the following command:

```
$ chmod +x IBM-Hyper-Scale-Manager_<version>.bin
```

5. Run:

```
./IBM-Hyper-Scale-Manager_<version>.bin
```

A menu similar to the following is displayed:

```
-----  
-----  IBM Hyper-Scale Manager <version>  -----  
-----  
IBM Hyper-Scale Manager version 5.0.2 is already installed on this server.  
Please select one of the following:  
1) Upgrade  
2) Uninstall  
3) Exit  
Please enter your choice > 2  
  
IBM Hyper-Scale Manager will be uninstalled.  
Are you sure you want to uninstall now? [y/n] >
```

6. From the menu, select the number that corresponds to Uninstall and press Enter.

7. Type y to run the uninstall.

IBM Hyper-Scale Manager will be uninstalled. The installation contains VMware integration data that will be lost during the uninstall.

Chapter 3. Configuration and administration

After installation and initial setup has been completed, you can perform various IBM Hyper-Scale Manager configuration and administration tasks as needed.

Migrating IBM Hyper-Scale Manager inventory from XIV Management Tools 4.x

If you have previously worked with IBM XIV Management Tools 4.x and have upgraded to IBM Hyper-Scale Manager 5.x, you can manually migrate your existing systems inventory to the newly installed IBM Hyper-Scale Manager.

About this task

Follow the steps to manually migrate your existing systems inventory from XIV Management Tools 4.x to the newly installed IBM Hyper-Scale Manager.

Procedure

1. Start by running the "first-time" login configuration to IBM Hyper-Scale Manager. See [“Completing initial setup”](#) on page 5.
2. Log in to the IBM Hyper-Scale Manager 5.x server via ssh and stop the HSM service with the `service xivms stop` command.
3. From the IBM Hyper-Scale Manager 4.x system, find the `systems.xml` file in the `$MSMS_HOME/server/data` directory.
4. Copy the `systems.xml` file to the target IBM Hyper-Scale Manager 5.x server and paste it in the same `$MSMS_HOME/server/data` directory (as on the 4.x server). Overwrite the existing `systems.xml` file.
5. Start the HSM service on the IBM Hyper-Scale Manager 5.x server with the `service xivms start` command.
6. Log in to the IBM Hyper-Scale Manager UI. The systems should be available in your IBM Hyper-Scale Manager inventory.

Changing the display language

While the default user interface display language is English, you can easily switch between an English or Japanese UI at any time.

About this task

Starting from IBM Hyper-Scale Manager system version 5.5.0, you can now see the entire UI in either English or Japanese. Complete the following steps to change your display language.

Important: The UI language configuration is not a personal user preference. Changing the setting changes it for all of the Hyper-Scale Manager users on this account.

Procedure

1. Once logged into the Hyper-Scale Manager, choose **Management Server > Manager Configuration** from the side bar.
2. From the Hub, click on **Language**.

3. From the **Set language** drop-down list, choose the desired display language (see [Figure 1](#) on page 10).

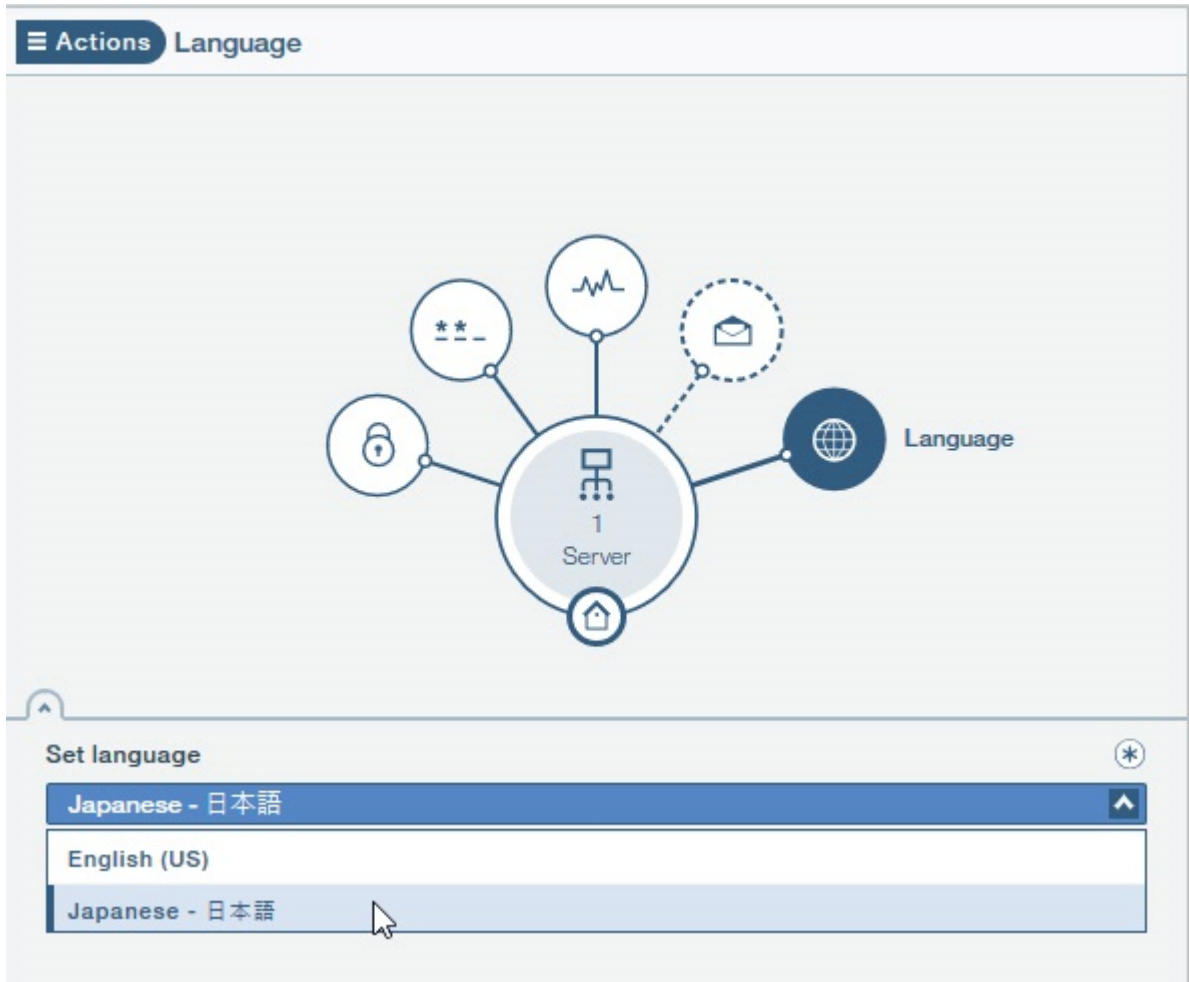


Figure 1. Changing the display language

4. Click **Apply** to continue.
The **Language Preferences** form appears.

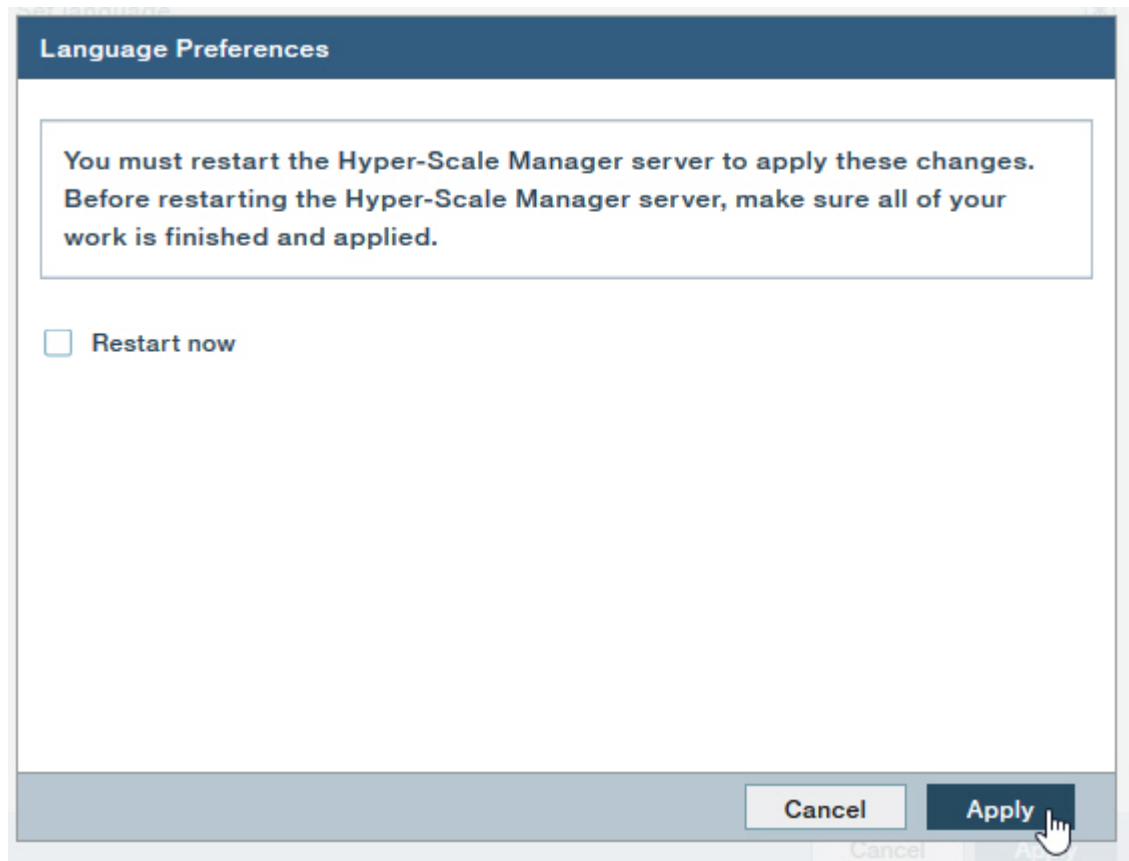


Figure 2. Language Preferences

5. In order to complete the **Set language** change, select the **Restart now** check box and click **Apply**.

Note: Make sure that all of your work is finished and changes are applied before restarting the Hyper-Scale Manager.

Language changes aren't applied to the display until the Hyper-Scale Manager is restarted.

If you only want to save your **Set language** changes but can only restart the Hyper-Scale Manager at a later time, leave the **Restart now** check box deselected and click **Apply**.

What to do next

To force a restart of the Hyper-Scale Manager, select **Restart > Restart Hyper-Scale Manager** from the **Actions** menu (see [Figure 3 on page 12](#)).

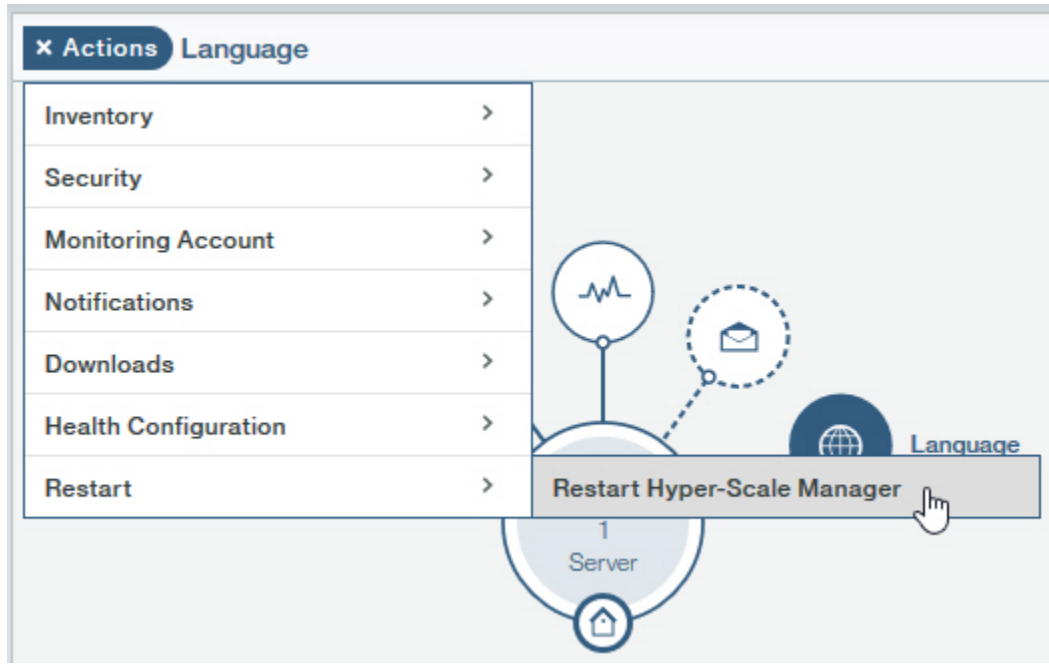


Figure 3. Restarting the Hyper-Scale Manager

Adding a system after initial setup

You can add additional systems to your inventory after initial setup.

About this task

Complete the following steps to add a system to your inventory.

Procedure

1. Log in to IBM Hyper-Scale Manager to further configure the Management Server.

Important: Regardless of what user you use to log in to Hyper-Scale Manager, the same username and password must be used on **all** of the systems in your inventory. If you change the password, note that the default time it takes for Hyper-Scale Manager to refresh the password is approximately 20 minutes. To update immediately, select **Actions > Inventory > Re-authenticate all Users**.

2. Choose **Management Server > Add systems to Monitor** from the side bar.
A new tab is opened in the table view and the **Add System** form is displayed:

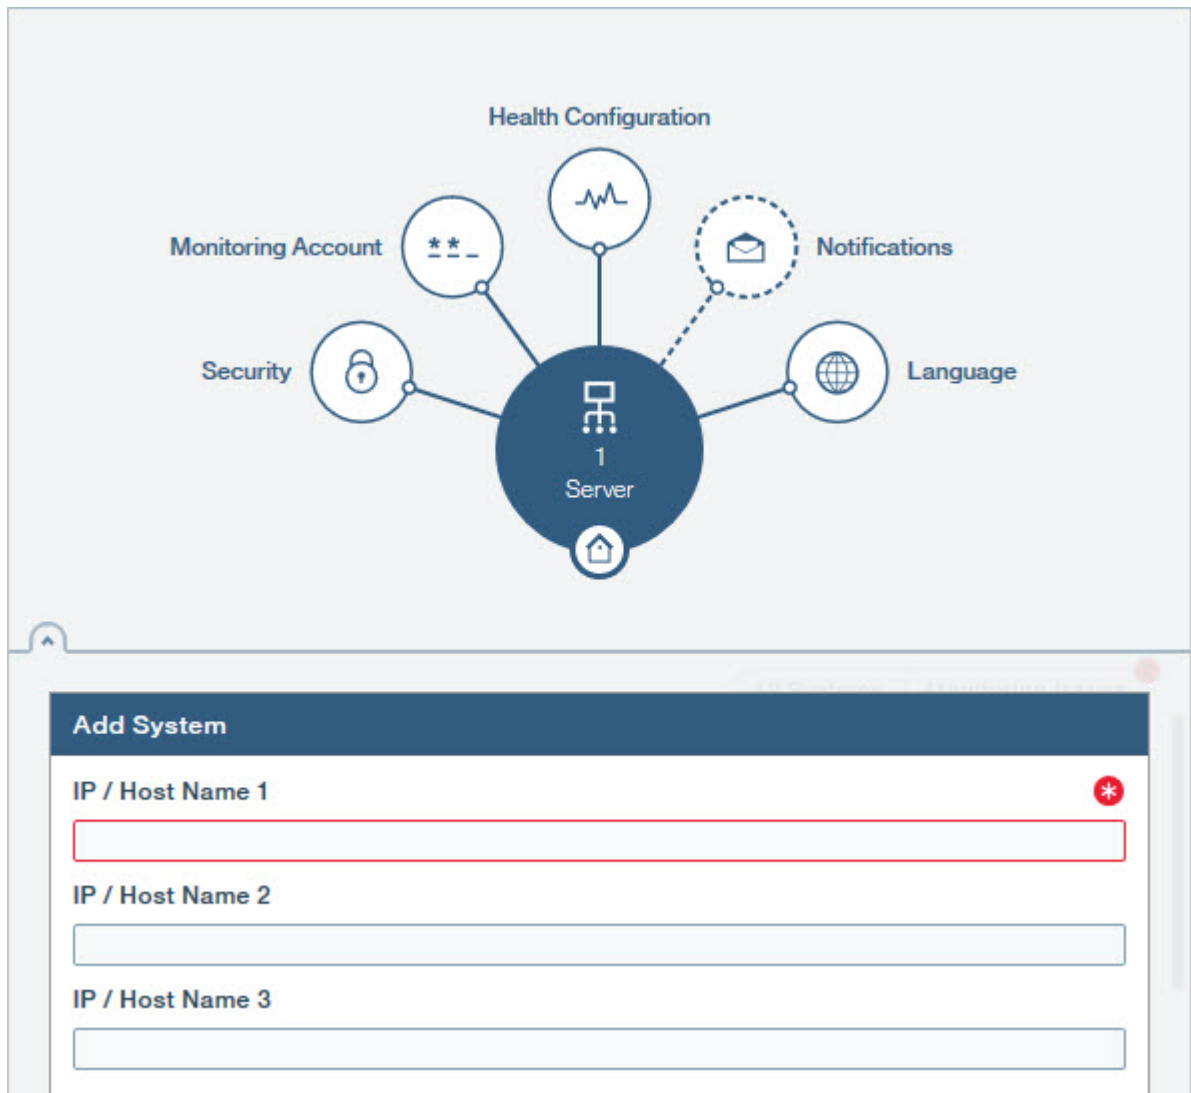


Figure 4. Adding a system to the inventory

3. Enter an IP address or Host Name for the system you want to add. Click **Apply**.
4. If a Monitoring Account doesn't exist on the newly added system, it will appear with an authentication error. From the **Actions** menu, select **Inventory** > **Fix all Authentication Errors**.

Note: You must be logged in to Hyper-Scale Manager as a predefined *admin* user to perform this operation.

5. (Optional) If you want to make changes to the monitoring account information used by the Management Server to monitor the systems in your inventory, click **Monitoring Account** on the Hub. To apply the changes to **all** of the monitored systems in your inventory, select **Actions** > **Monitoring Account** > **Set Account on Systems**.

Note: You must be logged in to Hyper-Scale Manager as a predefined *admin* user to perform this operation.

6. (Optional) Click **Notifications** on the Hub to allow and add users to receive notifications about Hyper-Scale Manager issues to the destination email.

☰ Actions Notifications

Subscribers (0)

1 Server

Why do I need notifications from IBM Hyper-Scale Manager?

SMTP configuration is utilized for enabling sending email messages to any recipient you configure in case of an IBM Hyper-Scale Manager service failure. The available SMTP email messages are: "Server is up", "Server is down", "Server has low disk space" and "System Machine Account is not authenticated"

Enabled *
 Disabled

SMTP Server Address * Port *

Sender Email Address *

Destination Email Address +

Email Address x

Email Address x

Figure 5. Adding a destination email

7. Click **Apply** to save the information entered. Click **Cancel** to exit without saving the information entered.

Configuring an access code

The access code restricts administrative configuration actions to the Management Server when you log in to IBM Hyper-Scale Manager.

About this task

Access to the Management Server is unrestricted by default, but you can add an access restriction at any time.

Procedure

1. If you want to add an **Access Code**, to the Management Server, click **Security** on the Hub.
2. Select **Restrict access**.
3. Type your access code and confirm it.
4. Click **Apply**.

When access to the server is restricted, the Hyper-Scale Manager login will appear in advanced mode.




Figure 6. Advanced login screen

5. To log in in **Advanced Mode**, enter the **System IP/DNS** address of one of your monitored systems, type your **Access Code**, and click **Advanced Login**.

Replacing the IBM Hyper-Scale Manager certificate

This section guides you through replacing the default IBM Hyper-Scale Manager certificate with your organization's certificate.

Before you begin

The IBM Hyper-Scale Manager is shipped with a default certificate. Use these steps to guide you through replacing the default certificate with your organization's certificate.

- Prepare a certificate file in PKCS#12 format.
- FTP the file to the hyperscale/files/upload folder.

Procedure

1. Log in to your Management Server and change the directory to the directory where you installed IBM Hyper-Scale Manager.
2. Type `./management_menu.sh`. The IBM Hyper-Scale Manager text-based user interface (TUI) User Menu is displayed:

```
-----  
----- IBM Hyper-Scale Manager v5.5.1 -----  
-----  
-----01/23/2019 05:06-----  
User Menu  
-----  
IBM Hyper-Scale Manager is running  
  
1) IBM Hyper-Scale Manager Inventory Options  
2) Show Audit Log  
3) Gather IBM Hyper-Scale Manager Logs  
4) Backup/Restore Options  
5) Change Manager Access Code  
6) Replace IBM Hyper-Scale Manager Certificate  
7) Change Network Settings  
8) SMTP Configuration  
9) Activate/Deactivate deduplication-aware capacity management  
10) Start/Stop Hyper-Scale Manager Service  
11) Exit  
Your Selection>
```

3. Choose Replace IBM Hyper-Scale Manager Certificate.

```
Your Selection>6  
  
This operation requires IBM Hyper-Scale Manager restart. Are you sure you want to  
continue? [Y/N] Y  
This script installs a new certificate.  
You need to prepare a certificate file in PKCS#12 format.  
Place the file to '/home/msms/hyperscale/files/upload' folder.  
  
Press any key to see the certificate files found under the  
'/home/msms/hyperscale/files/upload' folder
```

4. The available certificates that are in the *upload* folder are displayed on screen, similar to the following:

```
Certificate files found under the '/home/msms/hyperscale/files/upload' folder:  
-----  
Test-ErezCert_03_4096-cert.p12  
  
Please enter a certificate file name [] >
```

5. Enter a certificate file name and the key password.

```
Please enter a certificate file name [] >Test-ErezCert_03_4096-cert.p12  
Enter the certificate key password >  
The new certificate is successfully installed.  
The Hyper-Scale Manager service needs to be restarted so the changes will take effect.  
Restart now? [Y/N] >
```

6. Approve restarting the IBM Hyper-Scale Manager service.

Results

The server certificate is now replaced. You are returned to the User Menu.

Trusting the IBM Hyper-Scale Manager certificate

When you add a storage system to your inventory, you need to trust the IBM storage system certificate.

Before you begin

Managing certificates in this version of IBM Hyper-Scale Manager can be done using the CLI. For more information, refer to the *IBM FlashSystem A9000 Command-Line Interface (CLI) Reference Guide*.

Procedure

1. To view the status of the system certificate, authentication and license, select **Management Server > Manager Configuration**. The system inventory is displayed.
2. Click the **Actions** button in the upper right corner of one of the systems in the list to see the status of the certificate, authentication and license.

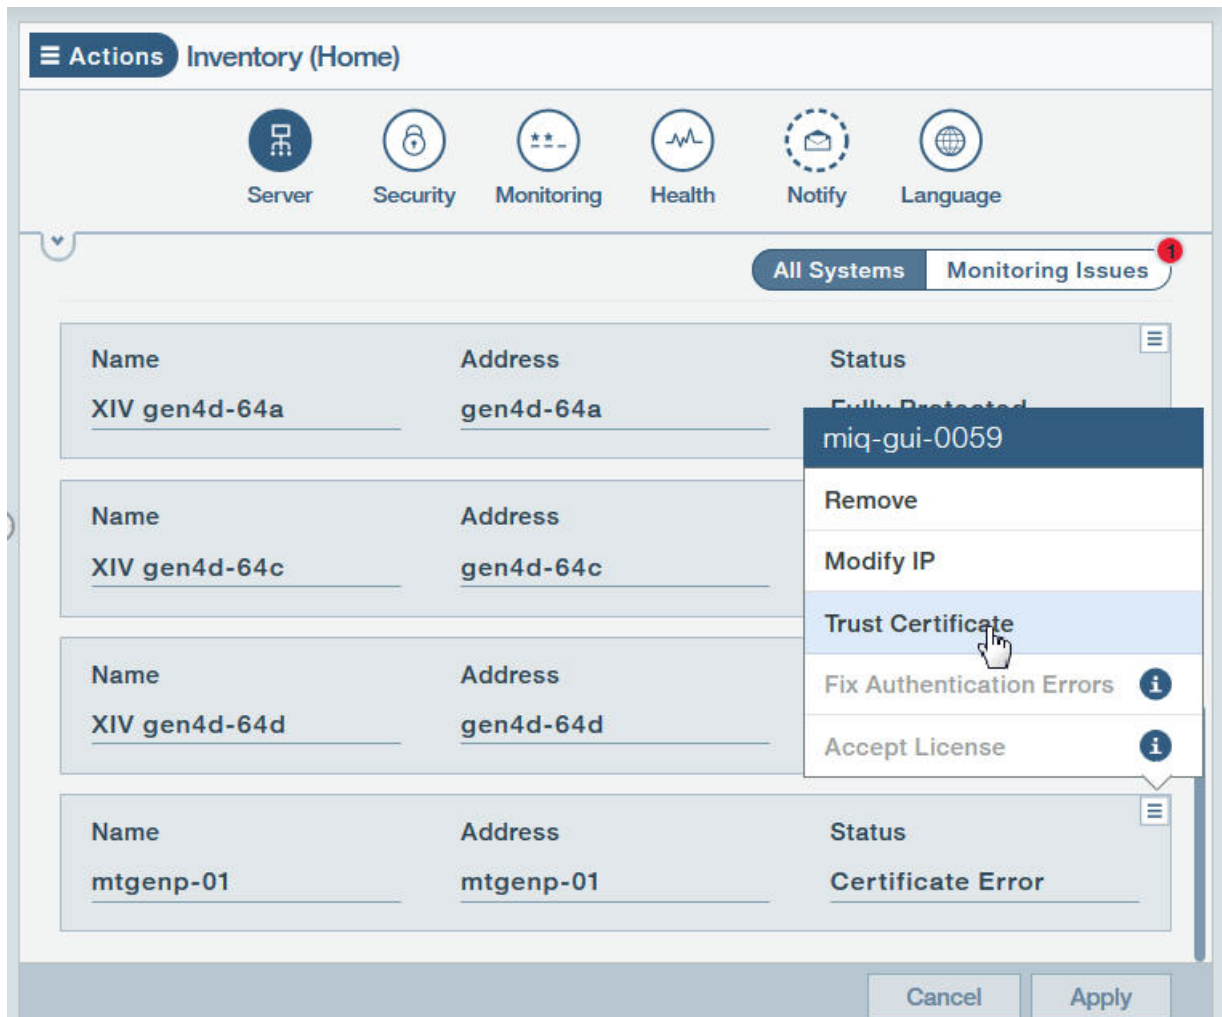


Figure 7. Viewing system certificate, licensing and authentication

3. If the system certificate or CA of the issuer is not trusted, then select **Trust Certificate**.
4. Click **Accept** to trust the certificate.

Backing up and restoring IBM Hyper-Scale Manager

The tasks in this section describe the best practice on how to back up and restore the inventory configuration of Hyper-Scale Manager. All of the backup and restore options are accessible from the *Management* menu of the text-based user interface (TUI).

Backing up IBM Hyper-Scale Manager

This section describes how to back up the inventory configuration of the IBM Hyper-Scale Manager.

About this task

The backup file is version-specific. It can be restored into a Hyper-Scale Manager of the same version. Whenever you upgrade Hyper-Scale Manager to a new version, create a new backup file.

The backup procedure backs up the following configurations:

- Hyper-Scale Manager inventory
- User settings, where information about users and the systems they work on is saved
- Network configuration
- System trusted certificate
- SMTP configuration

The backup process does not back up other configurations and data, such as:

- Network and NTP
- Monitoring Account
- Hyper-Scale Manager Access Code
- IBM Hyper-Scale Manager certificate
- IBM Hyper-Scale Manager logs

To configure these "not-to-be-saved" configurations, refer to the relevant sections of this user guide.

Procedure

1. Log in to the Hyper-Scale Manager using SSH.
2. Run `$MSMS_HOME/management_menu.sh`.
3. Select 4 on the management menu.
The Backup and Restore Menu is displayed on screen.

```
Backup and Restore Menu
-----
1) Backup
2) Restore
3) Schedule Daily Backup
4) Remove Scheduled Backup
5) Back to previous menu
Your Selection>
```

4. Select 1 to back up the database.

```
Your Selection>1
Backing up...
Backup completed successfully
Press any key to continue
```

The backup is stored in the `files/backup` directory under the `/hyperscale` directory in the installation path.

5. Press any key to return to the Backup and Restore Menu.

6. Export the backup to an external device.

What to do next

You can activate an automatic scheduled daily backup by selecting the `Schedule Daily Backup` option from the `Backup and Restore` Menu and specify the desired time. You can cancel it at any time by selecting `Remove Scheduled Backup`.

Restoring IBM Hyper-Scale Manager

This section describes how to restore the inventory configuration and preferences of IBM Hyper-Scale Manager from a backup.

About this task

The restore procedure restores the following configurations:

- Hyper-Scale Manager inventory
- User settings, where information about users and the systems they work on is saved
- Network configuration
- System trusted certificate
- SMTP configuration

The restore process does not restore other configurations and data, such as:

- Network and NTP
- Monitoring Account
- Hyper-Scale Manager Access Code
- IBM Hyper-Scale Manager certificate
- IBM Hyper-Scale Manager logs

To configure these not-restored configurations, refer to the relevant sections of this user guide.

Procedure

1. Copy the backup file to the `files/backup` folder in the `/hyperscale` directory in the installation path.
2. SSH using `root` to the Hyper-Scale Manager server.
3. Run `$MSMS_HOME/management_menu.sh`.
4. Select 4 on the management menu.
The `Backup and Restore` Menu is displayed on screen.
5. Select 2 to restore the database.

```
Backup and Restore Menu
```

```
-----  
1) Backup  
2) Restore  
3) Schedule Daily Backup  
4) Remove Scheduled Backup  
5) Back to previous menu  
Your Selection>
```

- a) Select a file (similar to the files below) to restore from:

```

Your Selection>2
Backup files to restore (the first is most recent):
-----
1) BACKUP_20160519_162219.tar.gz  3) BACKUP_20150519_162156.tar.gz
2) BACKUP_20160519_162211.tar.gz

Enter a number (from 1 to 3) to select a file or 'q' to quit >1

Are you sure that you want to restore BACKUP_20150519_162219.tar.gz? [Y/N] >y

```

- b) Click Y to confirm.
- c) Confirm again:

```

Restore DB is going to override your current configuration.
It is recommended to run a backup prior restoring.
Do you want to run backup of the current configuration before restoring? [Y/N] >y
Do you want to restore the existing Snapshot schedules [Y/N] >

```

Note: If you intend to restore your existing Snapshot schedules, and you have two or more active Hyper-Scale Managers, make sure that the Snapshot Schedule is configured on only one.

- d) Output similar to the following is displayed:

```

Backing up...
Backup completed successfully
Backup archive is in: /home/msms/hyperscale/files/backup/
BACKUP_20150519_162717.tar.gz
xivms v1.8.0.74 has stopped

Restoring from BACKUP_20160519_162219.tar.gz...
The restore is successfully completed.

Note: The following configurations were not restored ( the previous configurations
were kept ):
- System Machine Account password
- Manager Access Code
- Manager Certificate
If you wish to set them, please refer to the user guide for further explanations.

xivms v1.8.0.74 is starting...

```

- e) When the restore is complete, press any key to return to the Backup and Restore menu.
6. Go to [“Adding a system after initial setup”](#) on page 12 to configure your Monitoring Account information.

Managing connectivity

System connectivity is the connection between two systems that are monitored by IBM Hyper-Scale Manager, or the connection between one system that's monitored and one that's not monitored.

Defining connectivity allows you to connect your system to a remote system for purposes of mirroring or data migration.

Defining new connectivity

Before defining connectivity you must first assign a target system to which you want to connect.

About this task

Follow these steps to define new connectivity from IBM Hyper-Scale Manager.

Procedure

1. Select **System and Domains Views > System Connectivity**.
2. From the Hub, click on **Targets**.

The list of Targets is displayed in the **TARGETS** panel on the **System Connectivity** view. If there are no Targets connected to this System, the display will be empty.

3. At the top of the **TARGETS** panel, click the plus (+) sign to add a target.
The **Creating New Connectivity** form is displayed.

Status: Creating New Connectivity Define Connectivity

Monitored Remote System
 Unmonitored Remote System

Source System: XIV Gen4D-32-upperb

Remote System: XIV gen4d-64d

Source Domains: +

Remote Domains: +

Source Domain: /Global Space/

Remote Domain: /Global Space/

TARGET PARAMETERS

Migration
 Mirror / Mobility

Max Sync Job Rate (MB/s): 300

Max Resync Rate (MB/s): 300

Max Initialization Rate (MB/s): 100

Attached to Quorum Witness: None

Figure 8. Creating new connectivity

4. Add a **Remote System** that is either monitored by Management Server or unmonitored and define all of its parameters, including Multi-tenancy. Click **Apply**.
5. Click **Define Connectivity** to view the **System Connectivity** view, which displays the defined ports on the Source and Target systems.
A dotted line between the Source and Target systems shows that the connectivity is not yet defined.



Figure 9. System Connectivity view

Note: If the remote system target is Ethernet (iSCSI), you can create an IP Interface with VLAN for any of the ports showing on the **ETHERNET PORTS** form using the Add IP Interface form. Keep in mind, however, that:

- A VLAN can only be added during initial creation of an IP interface.
- System connectivity can only be implemented on a port with a single IP interface.

To access the Add IP Interface form, click the **Actions** menu in the upper-right corner of the row for the port that needs to be configured on the **ETHERNET PORTS** form and select **Add IP Interface**.

Add IP Interface

| | |
|----------------------------|------------------------|
| Port | Module |
| 2 | 1:Module:2 |
| IP Interface Name * | VLAN Tag * |
| | |
| IPv4 Address | Default Gateway |
| 0.0.0.0 | 0.0.0.0 |
| Netmask * | |
| | |

Cancel Apply

Figure 10. Add IP interface form

6. Click **Auto Connect**. The system sees all of the ports on the Target and connects the ports. You can switch between the Source and Target views on the **FC PORTS** or **ETHERNET PORTS** form on the **System Connectivity** view to see the ports and the port status defined on both.
7. To see a graphical representation of the connections between the ports of the Source and Target systems, click **Connectivity Details** on the **System Connectivity** view to display the **System Connectivity** details view.




Figure 11. System Connectivity details view

Note: IBM Hyper-Scale Manager uses the term *Module* when referring to XIV data and interface modules or FlashSystem A9000 and A9000R grid controllers.

Note: To download a more detailed report, click **CSV** in the lower-right corner of the System Connectivity details view. The file includes details about all of the ports per module in the system separated by Source and Target, including the port number, port type, port address, configured rate (Gbit), current rate (Gbit), status, and role.

If a VLAN was added to the definition of an IP interface for an Ethernet port, this information will also be included in the report.

Defining connectivity on unmonitored systems

If you're defining connectivity on systems that aren't monitored by Management Server, you need to go to the Target display of the System Connectivity view.

Procedure

1. Click **Add**.

The **FC PORTS** or **ETHERNET PORTS** form is displayed.

2. Type the IP Address or WWPN of the remote system, depending on the port type, and make sure the status is *active*.
3. Open the IBM Hyper-Scale Manager UI on the remote system to add the ports on the Target.
4. Click **Auto Connect**.

Chapter 4. Viewing the dashboard and system health score

The **Dashboard** of the IBM Hyper-Scale Manager UI collects information for analysis and provides an overview of your managed storage systems in real-time.

The **Dashboard** displays an aggregated view of your system health, your data reduction rate and savings, provisioned capacity and performance metrics.

The **Health** widget monitors all of your system core components. System Health monitoring not only detects and raises *alerts* for immediate problems, but gives you *tips* per feature group that expose potential issues which may ultimately affect your system health and performance.

All *tips* and *alerts* are clickable, making it easier to jump directly to the problematic object.

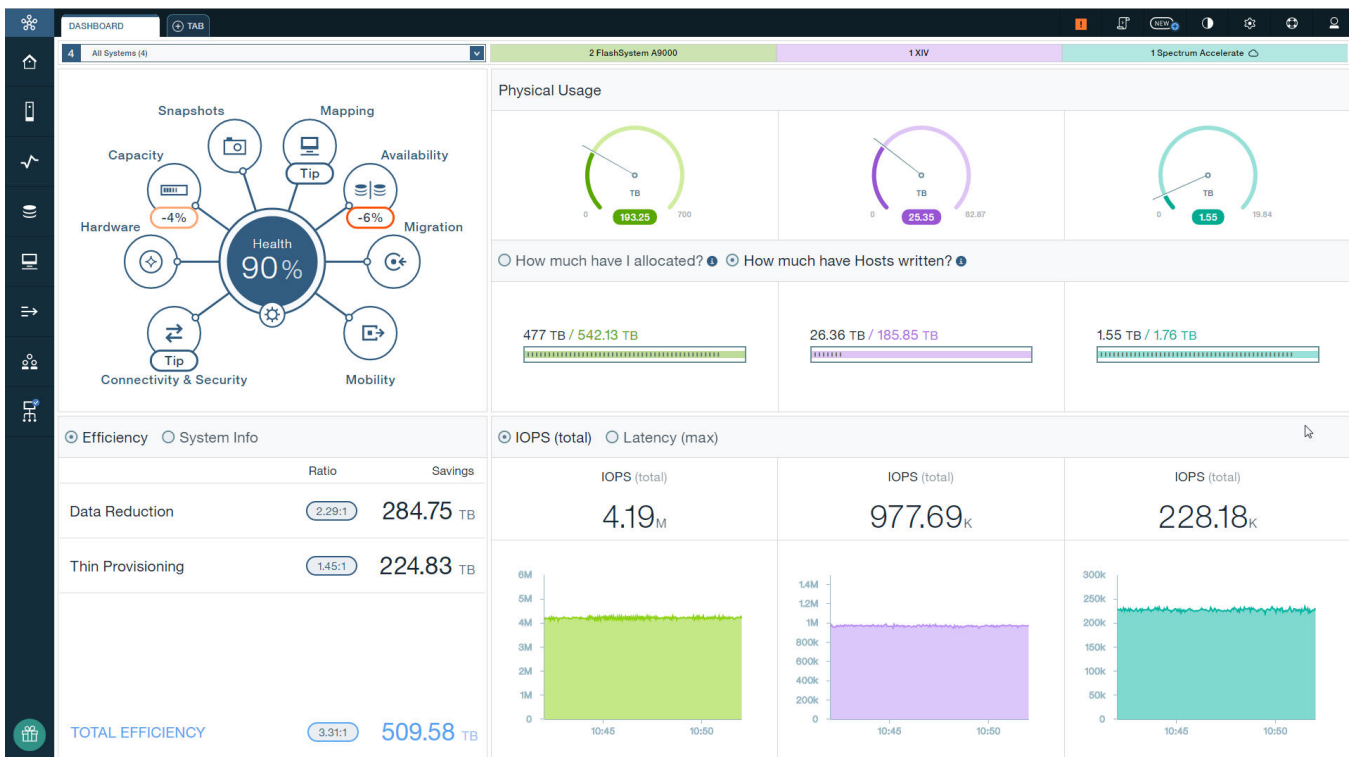


Figure 12. Health widget on the Dashboard

The total health status is dependent upon the system alerts and tips. Alerts reduce the overall health score percentage, while tips have no effect on the score.

Alerts appear for connectivity, hardware, and capacity and are broken down by severity: **Critical**, **Major**, **Minor**, or **Warning**.

Each severity has a calculation that contributes to the reduction of the overall health score. While the total health score can't be less than 10%, within each category the reduction of the health score is limited to -50%.

The default health reduction values appear in the [Table 1 on page 28](#). Click on the Settings icon on the Health widget to define how much your total dashboard health score will be reduced by each alert, per severity.

Table 1. Default health scores

| Severity | Health score |
|----------|--------------|
| Critical | -25 |
| Major | -6 |
| Minor | -3 |
| Warning | -1 |

Note: In some cases, the Health widget might display a score resulting from the failure of a top-level component. If this occurs, individual scores from alerts on internal components do not affect the score.

Tips appear within all of the feature groups on the Health widget, except hardware, and are further divided into different categories, such as:

- No Multipath Defined
- Host Connectivity
- Replication
- Snapshots

Here are some examples, where the tips would indicate some of the following type of System issues:

- Two systems are connected, but not defined as *multipath*, thereby leaving replication at risk, in case of a failure.
- Host connectivity where the Host is connected to one system, but not set up for *multipathing*, thereby leaving the Host at risk in case of failure.
- Volumes are not used and can be reclaimed.
- The mirroring or migration status of a single Volume, or multiple Volumes, indicates process errors.
- The Pool's Snapshots size is too large in ratio to the Pool size.

You can also toggle whether or not you want a tip to be displayed (see Figure 13 on page 29). This can be done from **Management Server > Manager Configuration** and clicking on **Health Configuration** on the Hub.

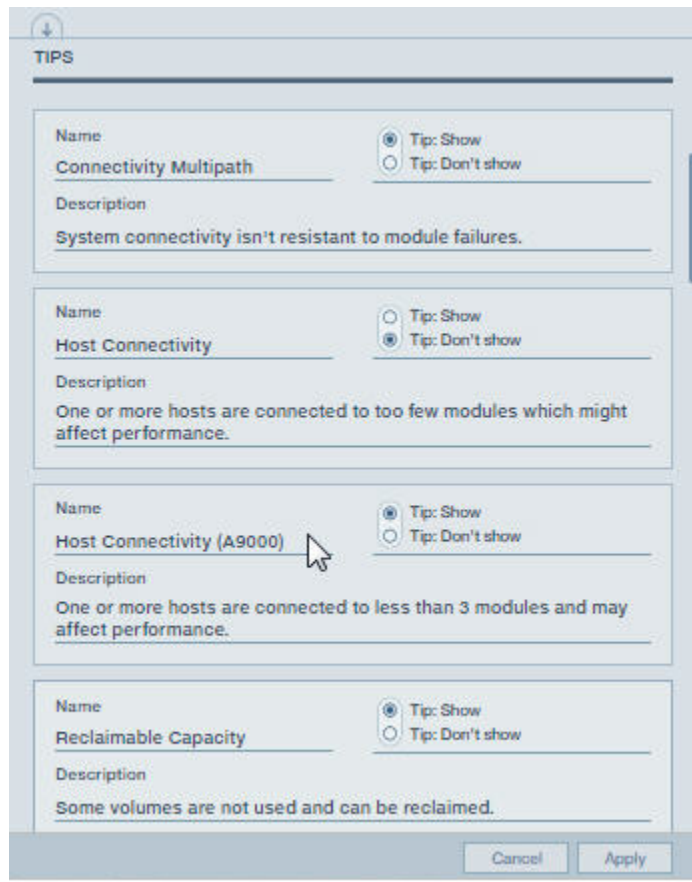


Figure 13. Configurable tips

Click on the alerts or tips to navigate to and fix the problematic objects so that you can improve your system's health score.

Chapter 5. Understanding and monitoring capacity

Storage vendors describe capacity in different ways, therefore interpreting capacity in your storage environment can be surprisingly tricky. In effort to simplify the notion of capacity in your storage environment, we've combined the concepts associated with capacity across all of the system types in your inventory.

The primary provisioned resource in storage systems is *capacity*. *Provisioning* is the process of associating resources with objects, such as: Systems, Domains, Pools (with their associated Snapshot space) and Volumes, as well as Hosts and Clusters.

Managing and utilizing the storage system's physical and projected capacities is a key aspect of system management. The [Figure 14 on page 31](#) illustrates capacity accounting, monitoring and planning.

On the right side of the figure, the **Physical size** (physical usable capacity) represents the amount of space that is physically available in the system to store data after data reduction. The **Stored** amount illustrates how much has been used from the **Physical size** of the system after data reduction. In addition, the physical thresholds set for the System indicate the current warning, minor, major and critical values.

With data reduction, more space becomes available for data storage. In [Figure 1](#), the **Projected system capacity** is the combined calculation of the reduction ratio you define for your expected workload (in this case, x5.0), together with the actual System reduction rate of your accumulated data workload. The **Allocated Volumes** value represents the total Volume size that you have created in the system. How much has been **Written by Hosts** represents the total data that was written by Hosts to all of the Volumes before it was reduced in the system.

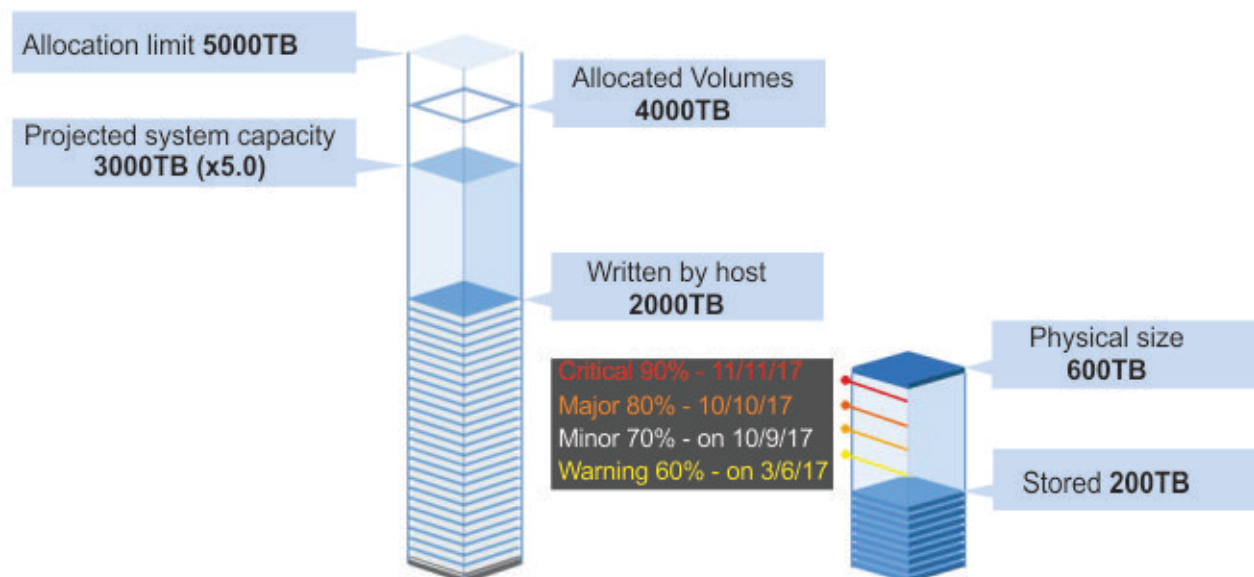


Figure 14. Projected vs. Allocated vs. Physical

Figure 14 on page 31 emphasizes the reduction rate of 1:10 while the expected reduction ratio is set to 1:5. This means that the reduction rate is better than the reduction ratio that was defined.

When you consider how data reduction will affect your space requirements, consider several factors:

- Think about the type of data that will be stored and how well that type of data typically can be reduced by using compression and data deduplication.
- Make basic projections by using typical compression rates for calculating potential savings. **Over allocation** can also save even more space.

For more information about thin provisioning, see the relevant Product Overview for your storage system:

- *IBM XIV Gen3 Product Overview* (GC27-3912). This guide can be found on the [IBM XIV Storage System Knowledge Center website](http://www.ibm.com/support/knowledgecenter/STJTAG) (www.ibm.com/support/knowledgecenter/STJTAG)
- *IBM FlashSystem A9000 Product Overview* (GC27-8583). This guide can be found on the [IBM FlashSystem A9000 Knowledge Center website](http://ibm.com/support/knowledgecenter/STJKMM) (ibm.com/support/knowledgecenter/STJKMM).
- *IBM FlashSystem A9000R Product Overview* (SC27-8558). This guide can be found on the [IBM FlashSystem A9000R Knowledge Center website](http://ibm.com/support/knowledgecenter/STJKN5) (ibm.com/support/knowledgecenter/STJKN5).

Following are some of the basic concepts associated with capacity in IBM Hyper-Scale Manager:

Projected System Capacity

The **Projected System Capacity** is the combined calculation of the **projected reduction ratio** you define with the accumulated system data, such as: growth rate and actual data reduction rate over time. The Projected System Capacity may fluctuate according to the continuous System calculation. Defining a fixed size value for the estimated reduction ratio keeps the Projected System Capacity calculation steady.

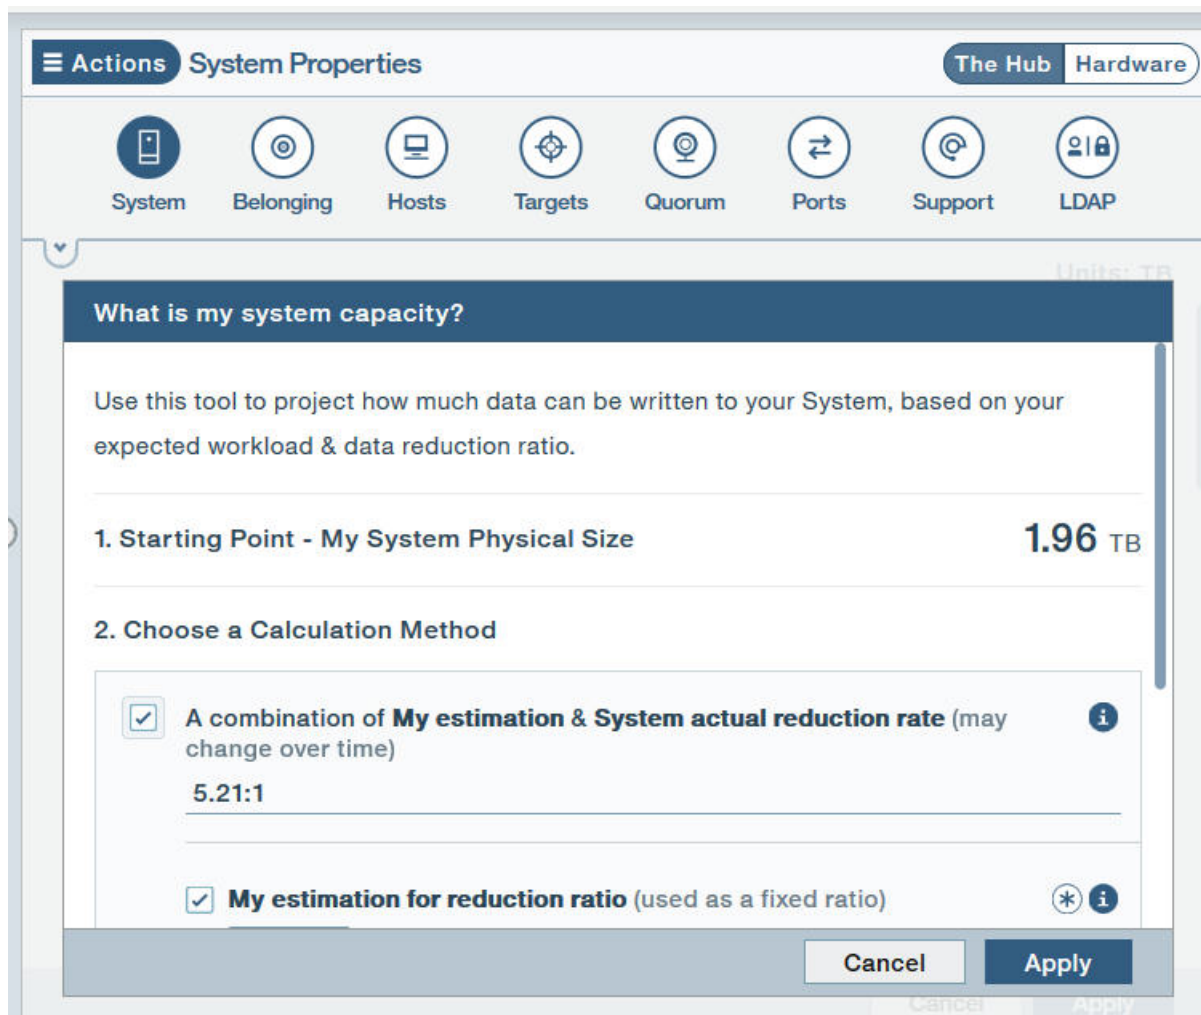


Figure 15. Projected System Capacity tool

Allocated Volumes & Snapshots

The total Volume and Snapshots size that you have created in the System.

Allocation limit

The default limit on the System available for Volume allocation.

Written (by Host)

Represents the total data written by hosts to the Volumes in the system before data reduction.

Physical size

The physical size (or physical usable capacity) represents the amount of space that is physically available in the system to store data after data reduction.

Physical usage

The actual total *physical size* used after data reduction.

Volumes in System

How much space has been reserved by creating Volumes and Snapshots out of the System's allocation limit.

Reclaimable capacity

The total allocation capacity of all Volumes which are not used in your environment; this is a total makeup of source Volumes (or their Snapshots) that aren't mapped and target Volumes that aren't mirrored. This capacity can, thus, be reclaimed and allocated to other hosts.

You can see the provisioned capacity across all of the System types in your inventory from the Hyper-Scale Manager Dashboard.

Monitoring your capacity

From the IBM Hyper-Scale Manager UI, you can monitor all of your systems' capacity from the dashboard across your "A" and "XIV" storage families at one time.

About this task

You can open additional tabs or use the dashboard drop-down menu to focus on System (tier) types or focus in on individual systems. Capacity can also be monitored from other various views, such as the System, Domains, Pools, Volumes, and Hosts views.

Procedure

1. From the dashboard, display **All Systems**. An overview of your managed storage systems in real-time is displayed, including any "off-premise" IBM Spectrum Accelerate systems.

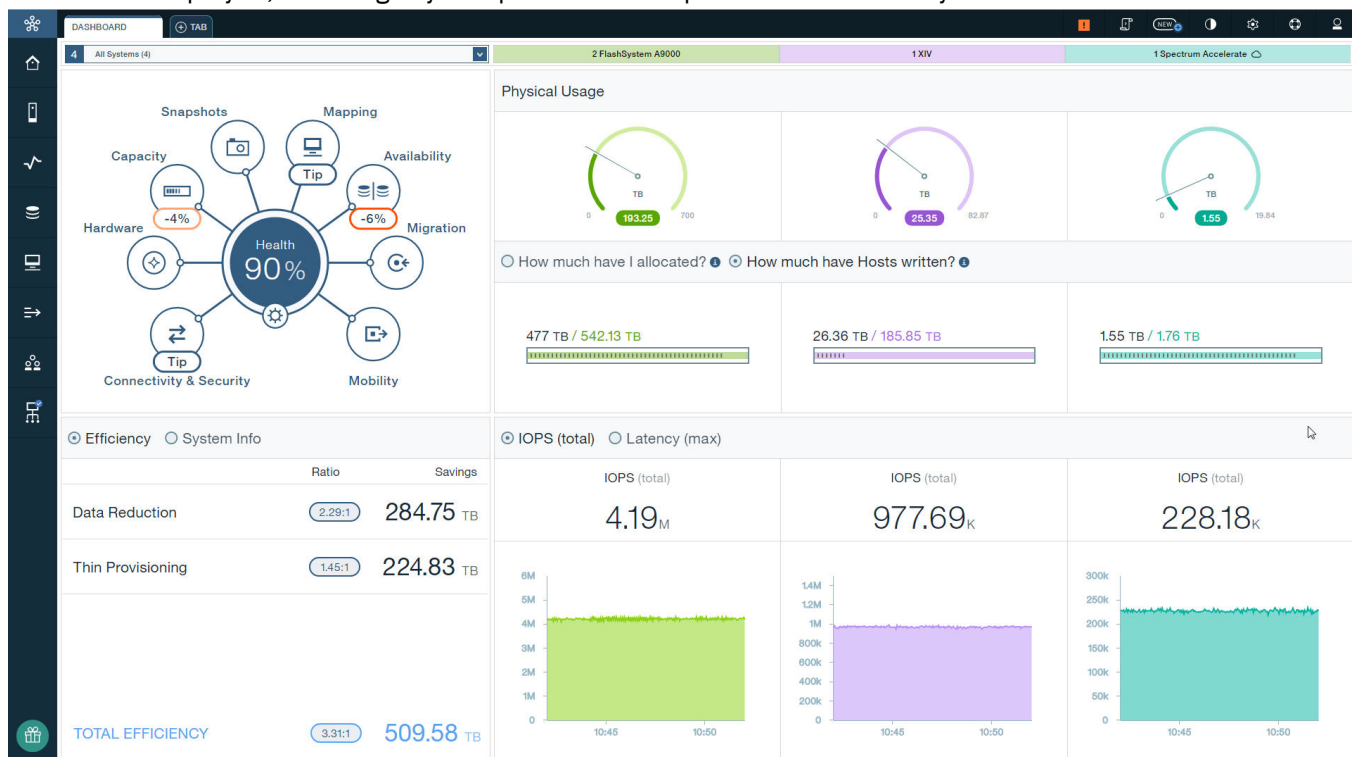


Figure 16. Dashboard

2. Click on one of the tabs above the gauges to display the selected tier, as shown in [Figure 17](#) on page 34.

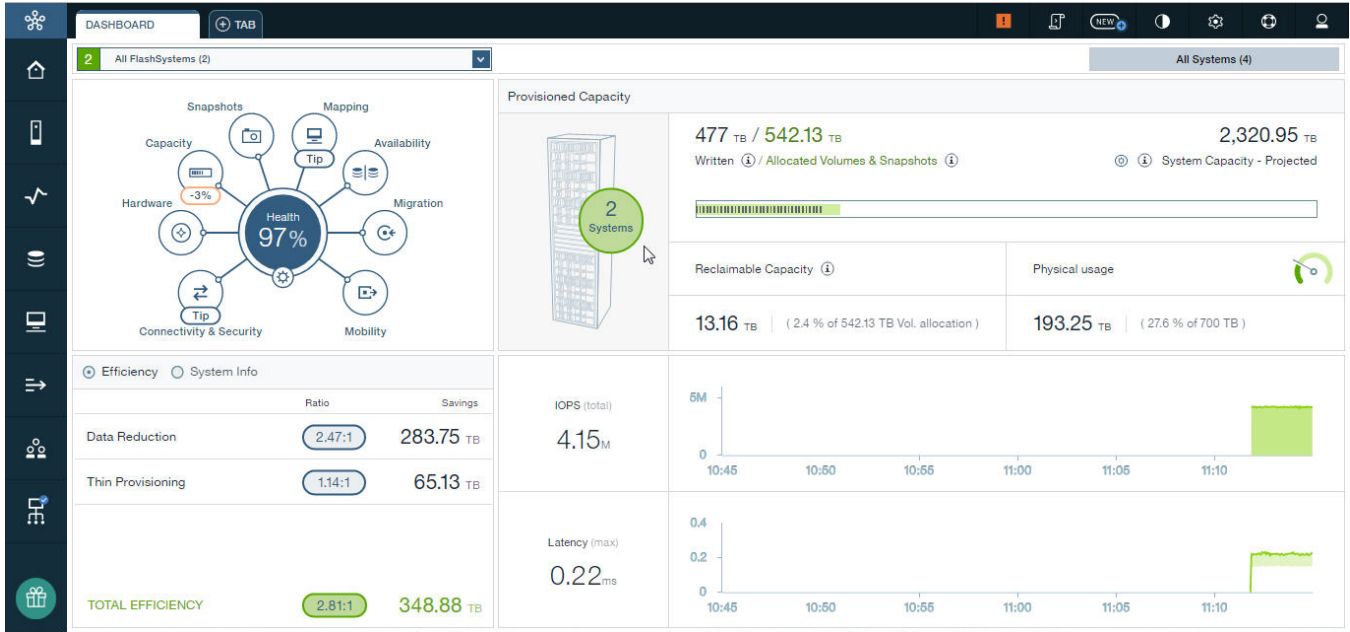


Figure 17. Select tier

3. Click on the Systems widget to display the System table and properties. Volume and Snapshot allocation, and usage before and after data reduction can be easily monitored using the tabs and gauges in the form.




Figure 18. Capacity view

In IBM FlashSystem A9000 and A9000R systems you can also see the physical capacity thresholds that trigger events about physical capacity consumption from the **Stored (after reduction)** gauge. An event of the defined severity is sent if the configured threshold value is reached or exceeded.

4. To configure the physical thresholds on an IBM FlashSystem A9000 or IBM FlashSystem A9000R system, select **Actions > Capacity > Configure Physical Thresholds**.

Chapter 6. Planning your capacity

IBM Hyper-Scale Manager provides capacity usage forecasts for IBM XIV Gen3, IBM Spectrum Accelerate, and IBM FlashSystem A9000 and A9000R systems.

Whether you prefer looking at Statistics graphs, or exporting and analyzing collected data in a table or PDF report, you can use the IBM Hyper-Scale Manager UI to collect the data usage forecast you need for planning your System capacity.

Collecting usage data for a storage system included in the inventory

The capacity usage for systems included in the inventory is collected once a day by the Management Server and stored in a **.dat** file on the server.

The capacity usage must meet several criteria for the forecast to be calculated:

- User must have sufficient access rights (role must be *Storage Administrator* or *Read-only*) for all of the systems included in the inventory.
- To present the trends, a sampling of 30 days is needed.
- A sufficient number of samples must be available. The forecast trend is not calculated if the number of samples is less than 30. If the system was not sampled for 14 days, then the system gathers 30 new samples, before the trend can again be presented.

If the sample fails, the Management Server samples the system every 15 minutes for a period of 12 hours. If the sample still fails, the operation is repeated the following day.

- System or storage pool utilization must be above 10%. If the system, or storage pool utilization, is less than 10%, no forecast is calculated.
- Trends cannot be calculated on pools that have no available space for volume allocation.
- If capacity is fluctuating or flat, or space utilization is decreasing, there is no trend.

Collecting usage data for a system that is removed from the inventory

The Management Server collects capacity data for systems that are listed on the inventory. Removing a system from the inventory implies stopping the data collection. However, to overcome situations in which the system was mistakenly removed from the inventory, or removed from the inventory for a short period of time, the Management Server applies the following rules on collecting capacity data for systems that are removed from the inventory:

- As long as the system is listed on the inventory, the Management Server collects and keeps its capacity data.
- Whenever the system is removed from the inventory, its capacity data is not immediately deleted. It is kept until the next time slot on which the data is collected from the machine.
- If the system is returned to the inventory prior to arriving to the next collection time slot, the capacity usage and its continuity are kept.
- If the system is removed from the inventory, it is impossible to reset its capacity data. To reset the capacity data, the system has to be listed in the inventory.
 - If the user chooses to reset capacity data for all systems, even non-monitored systems capacity data will be reset.

Generating a capacity planning report

Use this section to help you generate a capacity planning report from the IBM Hyper-Scale Manager UI.

About this task

The Capacity Planning report is generated from either the System, Domain, or Pool views.

The structure of the file name is:

- CAPACITY_PLANNING_REPORTS_YYYY-MM-DD_HH-MM-SS.zip

The zip contains the Capacity Planning PDF named:

- CAPACITY_PLANNING_REPORTS_YYYY-MM-DD_HH-MM-SS.pdf (The filename can be changed by the user.)

Procedure

Although there are multiple ways to generate a Capacity Planning Report, the following task describes how to generate a report from the Systems view.

Note: If you want to generate a Capacity report for multiple Systems, select the **Capacity Trend/Forecast Report** option from the **Actions** menu.

1. Select the system for which you want to generate a capacity report.
2. From the **Actions** menu, choose **Capacity > Capacity Trend/Forecast Report**.
A download bar appears at the bottom of your screen with a link to the containing folder of the zip file.
3. Open the zip folder and click to open the PDF.

Capacity planning report structure

The PDF file contains a formatted report with summarized informational graphics and graphs for each monitored System, Domain, and Pool.

The PDF report is divided into three parts:

- Part 1 - All System Types Usage (Physical)
- Part 2 - IBM FlashSystem A9000 and A9000R Pool Usage (Written)
- Part 3 - IBM XIV Gen3/Spectrum Accelerate Pool Usage (Physical)

Graphical representation illustrates the System and Pool capacity usage over periods of time, usage growth rates in GB per week, and detailed graphs per object capacity usage, respectively.

For example, [Figure 19 on page 39](#) shows an overall view for all of the systems in your inventory and a capacity forecast over the next 12 months:

Part 1 - All System Types Usage (Physical)

1.1 - System capacity usage over time

Over the next 12 months, out of 4 systems -




Figure 19. System allocation over time

The graph, [Figure 20 on page 39](#), ranks the system per used capacity growth rate.

1.2 - System by Usage Growth Rate

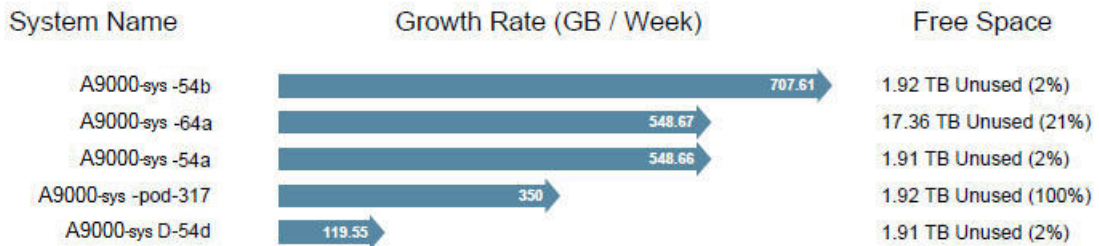


Figure 20. System by usage growth rate

Individual graphs for each system show the progression of the system capacity usage growth and a forecast for when the used capacity is expected to reach its defined thresholds. See [Figure 21 on page 40](#).

1.3 - System Usage - Detailed Graphs

— 80% Threshold — 90% Threshold — 100% Threshold

Sorted by the date when the used capacity is expected to hit the threshold

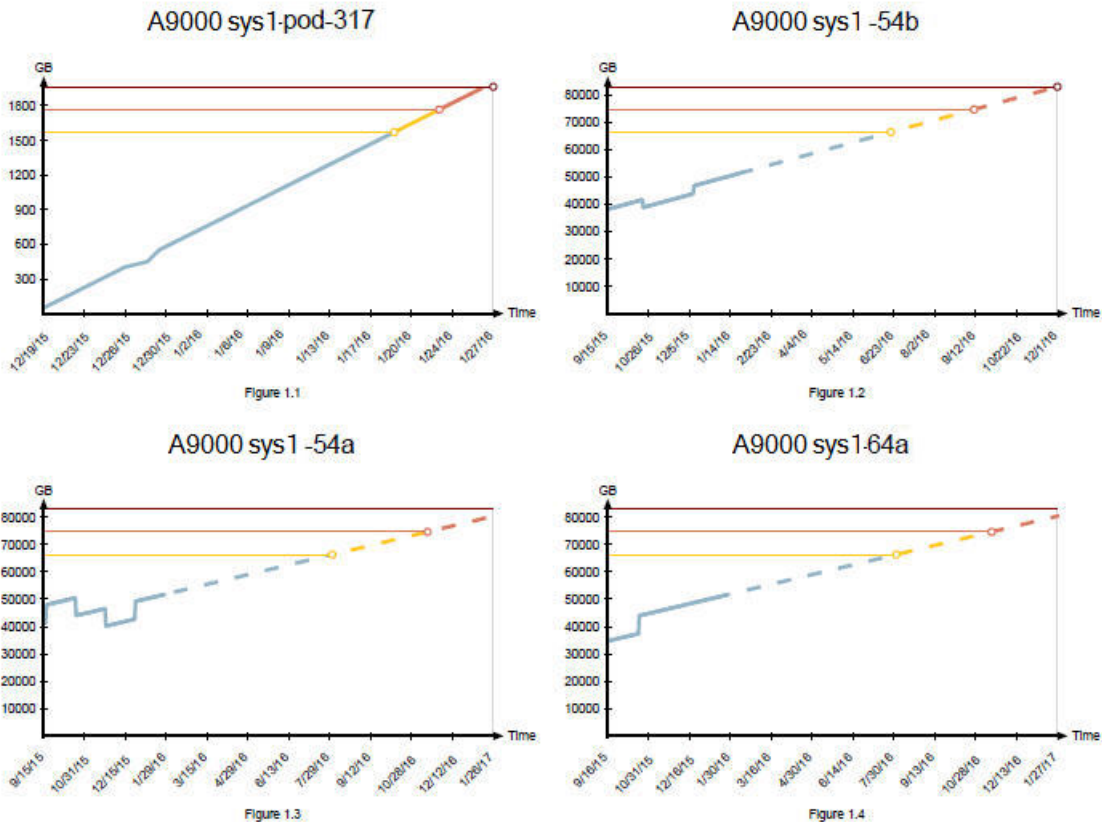


Figure 21. System usage - detailed graphs

Monitoring capacity statistics

Monitoring Capacity statistics helps you get a clearer grasp on your capacity usage patterns to create better and more established plans for your space utilization and future demands.

From the **Statistics Views** graphs in the IBM Hyper-Scale Manager UI, monitor your Capacity history and get a Capacity forecast for your Systems, Pools, and Hosts.

From the System view, choose from various filters to view the following System Capacity measurements:

Physical Usage

Shows how much data is used on your System out of your System's physical size.

Physical Usage + Forecast

Shows how much data is used on your System out of your System's physical size and what you can expect your System usage to be in the coming months.

Written/Allocated

Shows how much data has been written to your System out of the space allocated or reserved on your System for Volumes and Snapshots.

Written/Projected

Shows how much data has been written to your System from your System's Projected System Capacity after Data Reduction.

Written/Projected + Forecast

Shows how much data has been written to your System from your System's Projected System Capacity after data reduction. In addition, you can see the forecast of how much data you can expect to be written to your System in the coming months.

Projected/Actual

Shows your projected capacity (how much data you expect to be written to your System) and what's actually been written.

From the Pool view, choose from various filters to view the following Volume usage measurements:

Pool's Volume Usage

Shows how much total space has been used by Volumes out of the available space on the Pool.

Pool's Volume Usage + Forecast

Shows how much total space has been used by Volumes out of the available space on the Pool and how much you can expect to use in the coming months.

From the Host view, choose from various filters to view the following Host measurements:

Written by Host

Shows how much data has been written from the Host to your System.

Written by Host + Forecast

Shows how much data has been written from the Host to your System. The forecast shows what you can expect to be written from the Host in the coming months.

Understanding deduplication-aware capacity management calculations

Beginning with version 5.5.1, deduplication-aware capacity management calculations are available for IBM FlashSystem A9000 and A9000R systems running microcode version 12.3.1 and up.

Deduplication-aware capacity management is a deduplication enhancement that reports actual and more granular estimates of physical and reclaimable capacity per volume, as well as estimates of deduplication and compression ratios. It drills into vast amounts of metadata and analyzes it without any performance impact, using advanced heuristics that are proven to produce accurate estimates.

Using statistical data analysis, the deduplication-aware capacity management calculations compile the information of all the volumes on a system and calculate the *reclaimable* and *attributed* capacity per volume:

- The volume's *reclaimable* capacity is the portion of its physically stored data that is used only by that volume. When that volume is deleted, that capacity is assuredly reclaimed. When multiple volumes are deleted, the minimum assuredly reclaimed capacity is the sum of their unique capacities.
- The volume's *attributed* capacity is the amount of physical system capacity that can be fairly (proportionally) attributed to the volume, given that some of its data is shared with other volumes.

The deduplication-aware capacity management calculations are built to support future optimization tasks, such as determining the best system – from among multiple systems – to place a volume.

How effective are the calculations?

For the deduplication-aware capacity management calculations to provide accurate results, a sufficient amount of data and relative stability of the data during collection is required. Even if accuracy has suffered during a calculation cycle, however, the accuracy of the next calculation will be realigned.

The best calculation accuracy is achieved when:

- A volume has greater than 500 GB of data.
- A volume has a low change rate (i.e. less writes or deletes).
- Volumes have not been added or deleted recently.

System requirements

Ensure that your system meets all system requirements when using deduplication-aware capacity management calculations. For detailed information, refer to the IBM Hyper-Scale Manager release notes.

Note: If not all system requirements can be met for the deduplication-aware capacity management calculations, they can be deactivated post-installation or upgrade. They can also be reactivated at any time. For more information, see [“Deactivating IBM Hyper-Scale Manager deduplication-aware capacity management calculations”](#) on page 43.

Error handling

When alerts arise, error indications pertaining to deduplication-aware capacity management appear in the respective relevant columns of the table on the Volumes view.

If a problem occurs, the system will keep trying to collect data several times. When the alert is displayed in the table, the health widget on the dashboard will also indicate that the deduplication-aware capacity management calculation has failed and is not available.

Associated values available in the Volumes view

The following columns containing values accumulated from the deduplication-aware capacity management calculations may be displayed in the Volumes view if selected under **Columns > Add/Remove Columns:**

Stored (Attributed)

Shows the amount of the volume’s reclaimable capacity plus a proportional share of the volume’s physical capacity that is common (deduplicated) to additional volumes in the system.

Stored (Reclaimable)

Shows the amount of capacity that would be reclaimed if only that volume is deleted.

Reduction Savings

Shows the amount of data that was written to the system, minus the actual physical capacity that is needed to store that data past compression *and* deduplication. This is the *actual* saving.

Note: Compression and deduplication are always applied. The separation of the compression savings and deduplication savings columns in the RESTful API is only meant to realize their individual contribution to the overall reduction savings, as follows:

- **Compression Savings** would be equal to the Reduction Savings, if only compression could be applied.
 - **Deduplication Savings** would be equal to the Reduction Savings, if only deduplication could be applied.
-

Reduction Ratio

Shows the system-wide data reduction savings ratio. This value is used to project how much more data can be written to the system based on the estimated data reduction savings.

The **Stored (Attributed)**, **Stored (Reclaimable)**, and **Reduction Savings** values are also available to use as filters in the Volumes view.

Example use case

Suppose a capacity of 5 TB is needed for a new workload, but this exceeds the available (free) capacity in any of the systems in the environment. To vacate some space, your administrator considers deleting some volumes. Deleting volumes may be easy to plan without deduplication but is challenging to plan in a deduplicated environment. Use of deduplication-aware capacity management, however, can provide the solution.

Steps:

1. Display the volumes in the system and ensure that the **Stored (Attributed)** and **Stored (Reclaimable)** columns are visible in the table.
2. Filter the volumes with a **Stored (Reclaimable)** capacity > 100GB.
3. Sort the volumes in the table by the **Stored (Reclaimable)** column.

Note: The closer the reclaimable physical capacity is to the stored attributed value, the more the data on the volume is unique, and the better you can optimize your storage by vacating it. If the reclaimable capacity of a volume is low in comparison to its attributed stored data, you earn less by deleting it since you gain back a smaller percentage of your stored capacity.

4. Review the volumes in the table and select the volumes with equal or close to equal **Stored (Attributed)** and **Stored (Reclaimable)** values totaling approximately 5 TB.
5. Delete the selected volumes to free up the physical space you need for your workload.

Deactivating IBM Hyper-Scale Manager deduplication-aware capacity management calculations

If necessary, you can deactivate IBM Hyper-Scale Manager deduplication-aware capacity management calculations post installation or upgrade. Starting from version 5.5.2, these calculations are automatically activated during installation or upgrade of IBM Hyper-Scale Manager.

Before you begin

If your system meets all system requirements for use of IBM Hyper-Scale Manager deduplication-aware capacity management calculations, there is no need to deactivate them. Refer to the Hyper-Scale Manager release notes to determine if your system meets all requirements.

Procedure

Complete the following steps to deactivate IBM Hyper-Scale Manager deduplication-aware capacity management calculations in a system of version 5.5.1 or later.

1. Log in to your Management Server and change the directory to the directory where you installed IBM Hyper-Scale Manager.
2. Type `./management_menu.sh`.

The IBM Hyper-Scale Manager text-based user interface (TUI) User Menu is displayed:

```
-----
----- IBM Hyper-Scale Manager v5.5.3 -----
-----05/27/2021 05:06-----
User Menu
-----
IBM Hyper-Scale Manager is running

1) IBM Hyper-Scale Manager Inventory Options
2) Show Audit Log
3) Gather IBM Hyper-Scale Manager Logs
4) Backup/Restore Options
5) Change Manager Access Code
6) Replace IBM Hyper-Scale Manager Certificate
7) Change Network Settings
8) SMTP Configuration
9) Activate/Deactivate deduplication-aware capacity management
10) Start/Stop Hyper-Scale Manager Service
11) Exit
Your Selection>
```

3. Press 9.

The **Activate/Deactivate deduplication-aware capacity management** screen is displayed.

4. Select the appropriate menu options to complete the deactivation process.

The Management Server restarts, and IBM Hyper-Scale Manager deduplication-aware capacity management calculations are deactivated.

Note: The IBM Hyper-Scale Manager deduplication-aware capacity management calculations can be reactivated at any time, using the same steps above.

Moving capacity data between Management Server instances

The raw capacity planning data can be imported or exported from one IBM Hyper-Scale Manager to another and can also be reset to allow for collecting it anew.

Exporting the raw capacity data

The raw capacity planning data can be transferred from one IBM Hyper-Scale Manager to another.

About this task

The raw data that was collected on one Management Server can be used by another Management Server to maintain the continuity of system history.

Procedure

1. From the IBM Hyper-Scale Manager text user interface (TUI) main menu, select **IBM Hyper-Scale manager Inventory Options**.
2. From the **IBM Hyper-Scale manager Inventory Options** menu, select **Manage Capacity Planning Data**.
3. Select **Export Capacity Data** from the **Manage Capacity Planning Data** menu.

```
-----  
----- Hyper-Scale Manager v5.x -----  
-----  
Manage Capacity Planning Data  
-----  
1) Import Capacity Data  
2) Export Capacity Data  
3) Reset Capacity Data  
4) Back to previous menu  
Your Selection>2  
The capacity data file (*.dat) was exported to the (/home/msms/hyperscale/files/export)  
folder  
Press any key to continue
```

4. The file is exported. Press any key.

What to do next

Using the maintenance account, export the capacity data file created in this task from IBM Hyper-Scale Manager using SFTP to the target IBM Hyper-Scale Manager. The .dat file is stored in the export folder.

Importing the raw capacity data

The capacity planning raw data can be transferred from one IBM Hyper-Scale Manager to another.

Before you begin

Prepare a capacity data file that was created by another IBM Hyper-Scale Manager. SFTP to the target IBM Hyper-Scale Manager using the maintenance account. Put the .dat file in the upload folder.

About this task

A report that was created on one Management Server can be used by another Management Server to maintain the continuity of system history.

Procedure

1. From the IBM Hyper-Scale Manager text user interface (TUI) main menu, select **IBM Hyper-Scale manager Inventory Options**.
2. From the **IBM Hyper-Scale manager Inventory Options** menu, select **Manage Capacity Planning Data**.
3. Select **Import Capacity Data** from the **Manage Capacity Planning Data** menu.

```
-----  
----- Hyper-Scale Manager v5.x -----  
-----  
Manage Capacity Planning Data  
-----  
1) Import Capacity Data  
2) Export Capacity Data  
3) Reset Capacity Data  
4) Back to previous menu  
Your Selection>1  
Put the capacity data file (*.dat) in the (/home/msms/hyperscale/files/upload) folder  
Press any key to continue
```

4. Select from the available files in the upload folder. Press any key to continue. The file is imported. Restart the server.

What to do next

Whenever you generate a new report, the IBM Hyper-Scale Manager unifies the imported data according to the following continuity rules:

- Data of systems that are not managed by both IBM Hyper-Scale Manager instances is no longer tracked.
- Data for systems that were already tracked by both IBM Hyper-Scale Manager instances will be overwritten to avoid duplicates.
- Data for systems that are currently tracked and whose data was not imported remains unchanged.

Resetting the raw capacity data

The raw capacity planning data can be reset to allow for collecting it anew.

About this task

To clear the system history from irregularities (that is, machine re-purposing), you can clear the machine history that is collected by IBM Hyper-Scale Manager and start gathering data from scratch. You can reset the capacity data for a single system, or for all of the systems that are managed by IBM Hyper-Scale Manager.

Note: The system has to be tracked for its data to be reset.

Procedure

1. From the IBM Hyper-Scale Manager text user interface (TUI) main menu, select **IBM Hyper-Scale manager Inventory Options**.
2. From the **IBM Hyper-Scale manager Inventory Options** menu, select **Manage Capacity Planning Data**.
3. Select **Reset Capacity Data** from the **Manage Capacity Planning Data** menu.
In the following example, the capacity data for a system named *mn52* is reset.

```
-----  
----- Hyper-Scale Manager v5.x -----  
-----  
Manage Capacity Planning Data  
-----  
1) Import Capacity Data  
2) Export Capacity Data  
3) Reset Capacity Data  
4) Back to previous menu  
Your Selection>3  
Please choose which system(s) capacity data to delete:  
  system - system address to delete its capacity data  
  --all - delete all systems capacity data  
> mn52  
Are you sure you want to delete all capacity data for system: mn52? [Y/N] >y  
Capacity data was reseted for:  mn52  
Press any key to continue
```

4. Press any step to return to the **Manage Capacity Planning Data** menu.

Chapter 7. Troubleshooting latency issues

IBM FlashSystem A9000 and A9000R storage systems version 12.2.1 (or later) simplify root cause analysis of latency issues by breaking down latency measurements as follows.

- *Internal Latency* - The accumulative duration of the **internal system** processing and data transport of an I/O request.
- *External Latency* - The accumulative duration of processing and data transport of an I/O request **outside the system**.

Note: The external latency includes the portion of host latency which is inseparable by the system, since the system cannot record the time on the host ports.

- *Total Latency* - The sum of both the internal latency and external latency.

The following illustration differentiates between internal and external latency measurements in a simplified manner. An actual timing chart would perhaps be more detailed, have differences for read and write operations, and may include repetitions, but all are beyond the scope of this document.

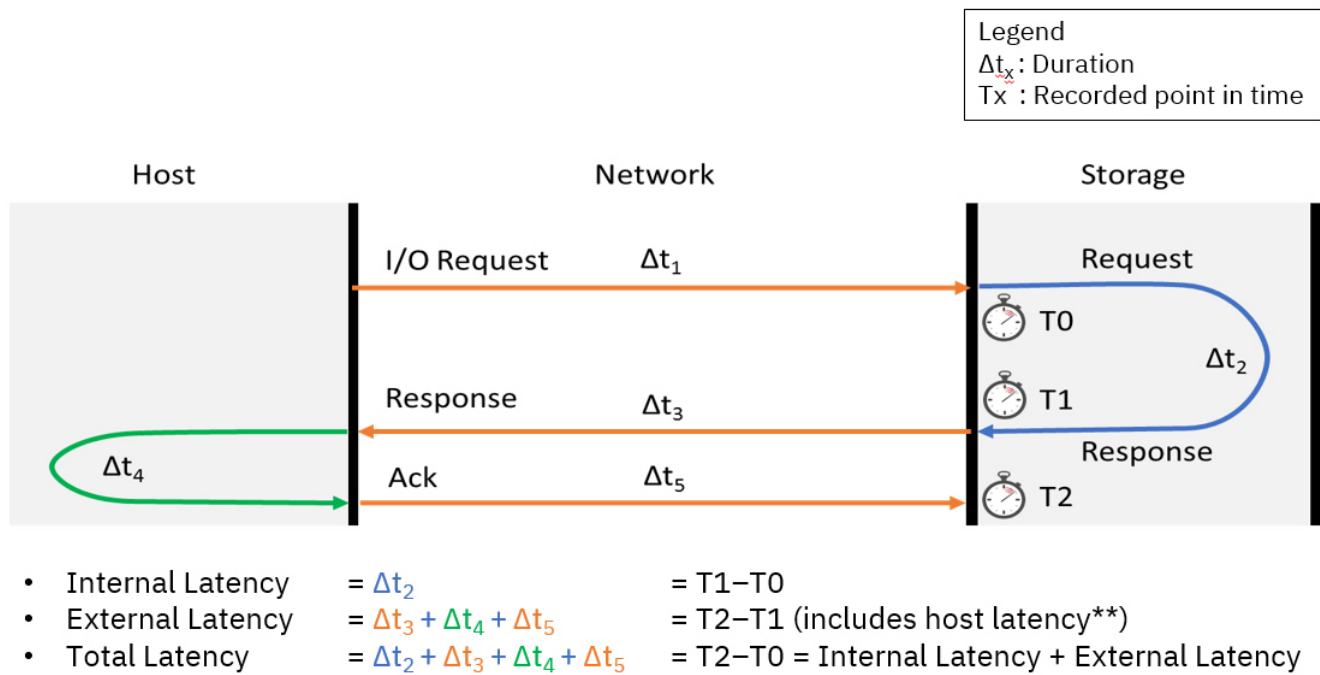


Figure 22. Illustration of the latency breakdown

A potential fault leading to a latency issue can be easily narrowed down to "inside the system" or "outside the system" by viewing your statistics in IBM Hyper-Scale Manager.

Latency can be viewed from either the System Dashboard, or the System and Host Statistics views. For example, from the **Hosts** view, select a single Host and choose **Current Latency** or **Historic Latency** statistics.

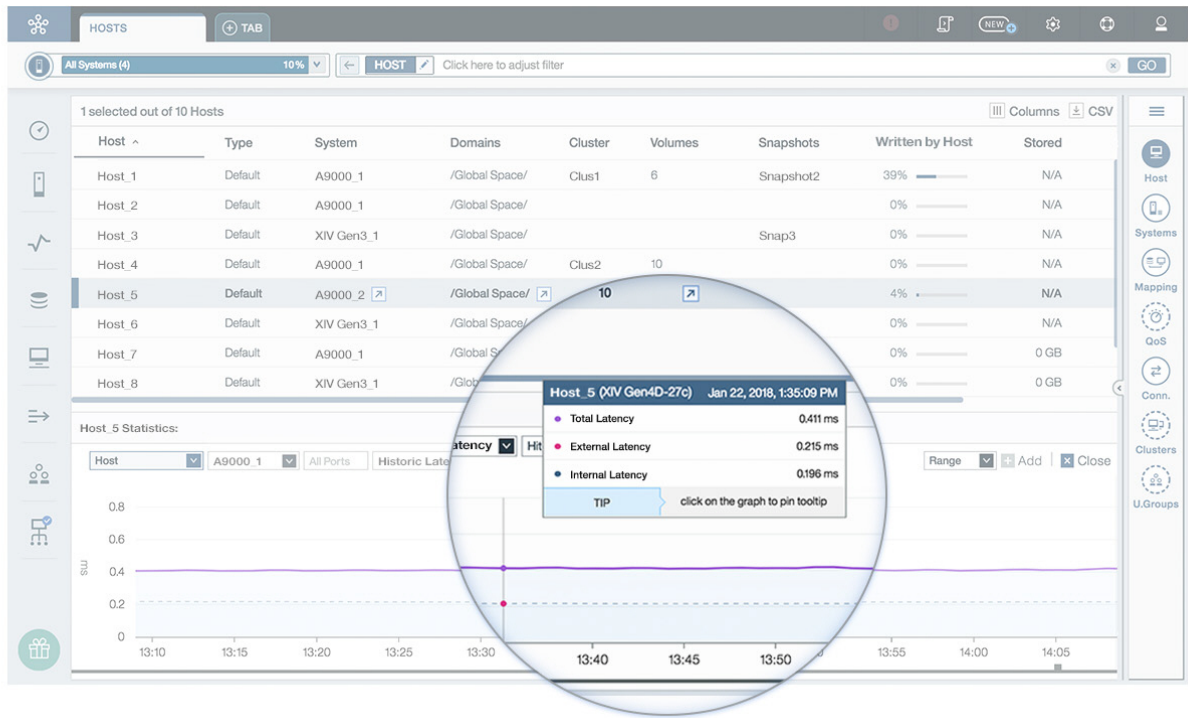


Figure 23. Latency view from Hyper-Scale Manager

Hover over the graphs to pinpoint any latency-related issues to better understand what's causing your latency.

Notices

These legal notices pertain to the information in this IBM Storage product documentation.

This information was developed for products and services offered in the US. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
USA*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119*

Armonk, NY 10504-1785
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at [Copyright and trademark information website \(www.ibm.com/legal/copytrade.shtml\)](http://www.ibm.com/legal/copytrade.shtml).

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

Index

Numerics

80%/90%/100% threshold [38](#)

A

Accounting

Capacity management [31](#)

Understanding capacity [31](#)

Adding Systems [12](#)

administration tasks [9](#)

analytics [44](#)

B

Backing up IBM Hyper-Scale Manager [19](#)

backup [19](#)

backup folder [44](#)

C

capacity analytics [44](#), [45](#)

Capacity Management

Provisioning capacity

Accounting [31](#)

Effective capacity [31](#)

Physical capacity [31](#)

Reclaimable capacity [31](#)

Total allocation [31](#)

Written by Hosts [31](#)

capacity planning [38](#)

Capacity Planning PDF Report [38](#)

Capacity report

generating a report [37](#)

certificate

of the IBM Hyper-Scale Manager [16](#), [18](#)

change

language [9](#)

comments, sending [viii](#)

configuration tasks [9](#)

Connectivity

Defining new connectivity [21](#)

managing [21](#), [47](#)

D

Dashboard

health score [27](#)

deduplication-aware capacity management

deactivating and reactivating [43](#)

Deduplication-aware capacity management [41](#)

display [9](#)

Domain usage [38](#)

E

edition notice [2](#)

Exporting capacity data [44](#)

F

feedback, sending [viii](#)

forecast [38](#)

forecasted capacity [38](#)

G

Generating a capacity planning report [38](#)

H

how to back up [19](#)

Hyper-Scale Manager

Completing initial setup [5](#)

I

IBM Hyper-Scale Manager

Management server [3](#)

Migrating your inventory from XIV Management Tools
4.x [9](#)

User Interface (UI) [3](#)

IBM Hyper-Scale Manager certificate [16](#), [18](#)

Importing capacity data [44](#)

Installation

IBM Hyper-Scale Manager [3](#)

Management server [3](#)

Introduction [1](#)

inventory configuration and preferences [20](#)

L

language

change [9](#)

Logging in

Advanced login [12](#)

M

maintenance account [44](#)

Management server

Installation [3](#)

Management Server

Monitoring account [3](#)

Migrating to 5.x [9](#)

Monitoring account

Management Server [3](#)

Monitoring and statistics [40](#)

Monitoring capacity [33](#)

moving capacity data [44](#)

N

notices
 legal [2](#)

P

PDF file [38](#)
Pool usage [38](#)
pools statistics [38](#)

R

resetting the raw capacity data [45](#)
restart
 IBM Hyper-Scale Manager [9](#)
restore [19](#)
Restoring IBM Hyper-Scale Manager [20](#)

S

score [27](#)
SFTP [44](#)
Standalone application [44](#)
structure of the PDF file [38](#)
System allocation [38](#)
System capacity
 planning [37](#)
System connectivity [21](#)
System Health [27](#)

T

The report legend [38](#)
threshold [38](#)

U

Uninstalling
 IBM Hyper-Scale Manager [7](#)
Upgrading
 IBM Hyper-Scale Manager [5](#)
upload folder [44](#)

V

Virtual appliance [44](#)



Printed in USA