

IBM Hyper-Scale Manager
Version 5.1.1

User Guide



Note

Before using this information and the product it supports, read the information in Notices.

Edition Notice

This edition applies to IBM Hyper-Scale Manager version 5.1.1 and to all subsequent releases and modifications, until otherwise indicated in new editions.

© Copyright IBM Corporation 2013, 2016.

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
Sending your comments	viii
Chapter 1. Introduction	1
Chapter 2. Installing and configuring IBM Hyper-Scale Manager	3
Installing IBM Hyper-Scale Manager	3
Setting up IBM Hyper-Scale Manager	5
Adding a system after initial setup	6
Replacing the IBM Hyper-Scale Manager certificate	9
Trusting the IBM Hyper-Scale Manager certificate	10
Upgrading IBM Hyper-Scale Manager version 5.x	11
Migrating IBM Hyper-Scale Manager inventory from version 4.x	12
Chapter 3. Backing up and restoring IBM Hyper-Scale Manager	13
Backing up IBM Hyper-Scale Manager	13
Restoring IBM Hyper-Scale Manager	14
Chapter 4. Viewing the dashboard and system health score	17
Chapter 5. Understanding and monitoring capacity	19
Monitoring your capacity.	21
Chapter 6. Planning your capacity	23
Generating a capacity planning report	24
Capacity planning report structure	24
Monitoring capacity statistics	26
Moving capacity data between Management Server instances.	27
Exporting the raw capacity data	27
Importing the raw capacity data	28
Resetting the raw capacity data.	29
Chapter 7. Managing connectivity	31
Defining new connectivity	31
Defining connectivity on unmonitored systems	35
Notices	37
Trademarks	38
Index	39

Figures

1. Advanced login screen	7
2. Adding a system to the inventory	8
3. Viewing system certificate, licensing and authentication.	10
4. Health widget on the Dashboard	17
5. Projected vs. Allocated vs. Physical	19
6. Select tier	21
7. Capacity view.	22
8. System capacity usage over time	25
9. System by usage growth rate.	25
10. System usage - detailed graphs	26
11. Creating new connectivity	32
12. Connectivity patch panel	33
13. Connectivity details.	34

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

You can find additional information and publications related to IBM Hyper-Scale Manager on the following information sources.

- IBM FlashSystem® A9000 on IBM Knowledge Center
- IBM FlashSystem A9000R on IBM Knowledge Center
- IBM XIV® Storage System on IBM Knowledge Center
- IBM Spectrum Accelerate™ on IBM Knowledge Center
- IBM Storage Redbooks® website (redbooks.ibm.com/portals/storage)
- IBM Flash Storage and Solutions marketing website (ibm.com®/systems/storage/flash)
- IBM XIV Storage System marketing website (ibm.com/systems/storage/disk/xiv)
- IBM Spectrum Accelerate marketing portal (ibm.com/systems/storage/spectrum/accelerate)

IBM Hyper-Scale Manager publications	Description
Release Notes	Describes what's new, requirements, compatibility, change log, and known issues information in the latest version.
User Guide	Describes how to install, configure, and operate IBM Hyper-Scale Manager.
Quick-Start Guide	Walks you through installing the IBM Hyper-Scale Manager software and initial configuration of the GUI.
REST API Specifications	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)
- IBM Support Portal website (www.ibm.com/storage/support)
- IBM Directory of Worldwide Contacts website (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/) offers customized search functions to help you find the publications that you need. You can view or download publications at no charge.

Sending your comments

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

Procedure

To submit any comments about this guide:

- Go to IBM Knowledge Center (ibm.com/support/knowledgecenter), drill down to the relevant product page, and then click the **Feedback** link that is located at the bottom of the page.

By adding a comment, you accept our [IBM Knowledge Center Terms of Use](#). Your comments entered on this IBM Knowledge Center site do not represent the views or opinions of IBM. IBM, in its sole discretion, reserves the right to remove any comments from this site. IBM is not responsible for, and does not validate or confirm, the correctness or accuracy of any comments you post. IBM does not endorse any of your comments. All IBM comments are provided "AS IS" and are not warranted by IBM in any way.

Comments (0) [Add Comment](#)

No Comments

Contact Privacy Terms of use Accessibility Feedback

Feedback

The feedback form is displayed and you can use it to enter and submit your comments privately.

- You can post a public comment on the Knowledge Center page that you are viewing, by clicking **Add Comment**. For this option, you must first log in to IBM Knowledge Center with your IBM ID.
- You can send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:
 - Exact publication title and product version
 - Publication form number (for example: SC01-0001-01)
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Note: When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Chapter 1. Introduction

IBM Hyper-Scale Manager is the official "A-line" hybrid cloud management application, supporting all-flash, disk and software-defined storage (SDS) systems. IBM Hyper-Scale Manager allows for unified management across the IBM FlashSystem A9000 and A9000R, IBM XIV Gen3 and IBM Spectrum Accelerate (on/off premises) family.

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.

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 Chapter 3, "Backing up and restoring IBM Hyper-Scale Manager," on page 13 for more information.)

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

- Using Virtualization Technologies

Virtualization technologies (for example, VMware tools) can be used to achieve High Availability.

From version 5.0.x, IBM Hyper-Scale Manager includes a new component, Cloud Integration, which is based on IBM Spectrum Control™ Base Edition. IBM Spectrum Control Base Edition is used to integrate storage systems with VMware systems and APIs, such as vSphere Web Client, VASA, vRealize Operations Manager, and vRealize Orchestrator. 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 IBM Hyper-Scale Manager UI login page. For more information about IBM Spectrum Control Base Edition, see IBM Spectrum Control Base Edition on the Knowledge Center.

Programmable automation and control of your storage systems can be implemented through the IBM Hyper-Scale Manager Representational State Transfer (REST) application programming interfaces (APIs).

IBM Hyper-Scale Manager provides a break-through 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.

Chapter 2. Installing and configuring IBM Hyper-Scale Manager

Installation and configuration 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.

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

Installing IBM Hyper-Scale Manager

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

About this task

Follow the steps to install Hyper-Scale Manager.

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 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 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
```

6. 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. >
```

Accept the installation location, or type a different path to change the location and 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] >
```

Type `y` to allow access. Traversal permissions (`x`-permission) is added to the selected directory. 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 above is also displayed.

7. Enter the communication ports at the prompts and press Enter.

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

Note: Since Cloud Integration is deployed as part of the IBM Hyper-Scale Manager package, a supported storage system can be added or removed only from the IBM Hyper-Scale Manager UI. When a system is added, it is displayed in the Cloud Integration interface from which you can run your usual Spectrum Control Base Edition operations. See http://www.ibm.com/support/knowledgecenter/STJKMM_12.0.2/hsm_rn_whatsnew.html.

Note: When you are choosing the Web UI ports, make sure not to use unsafe ports for Google Chrome or Mozilla Firefox. A complete list can be found at: [Which ports are considered unsafe on Chrome or Ports blocked by default in Mozilla.](#)

8. Enter `y` to begin the installation.

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

9. The license agreement is displayed. Press `1` and 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 IBM Confidential.

What to do next

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

- Outbound:
 - SMTP - 25
 - Communication with IBM FlashSystem A9000 and A9000R systems – 7778
- 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 iptables 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 Setting up IBM Hyper-Scale Manager.

Setting up IBM Hyper-Scale Manager

Run the "first-time" configuration to set up the first system in your storage environment.

About this task

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 Management Server. Make sure to use the same user credentials for the monitoring account on all of your systems.
 - 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.

- If you are using an LDAP domain, enter the name and password that exists in the LDAP server. Select the check box for UPN notation if your LDAP server requires a domain (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.

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 Hyper-Scale Manager to further configure the Management Server using either the default username `admin` and password `adminadmin` or the username and password that you defined. By default, access to the Management Server is unrestricted, but you can add an access restriction at any point (see step 4 on page 8). When advanced access is restricted, the login will appear like this:

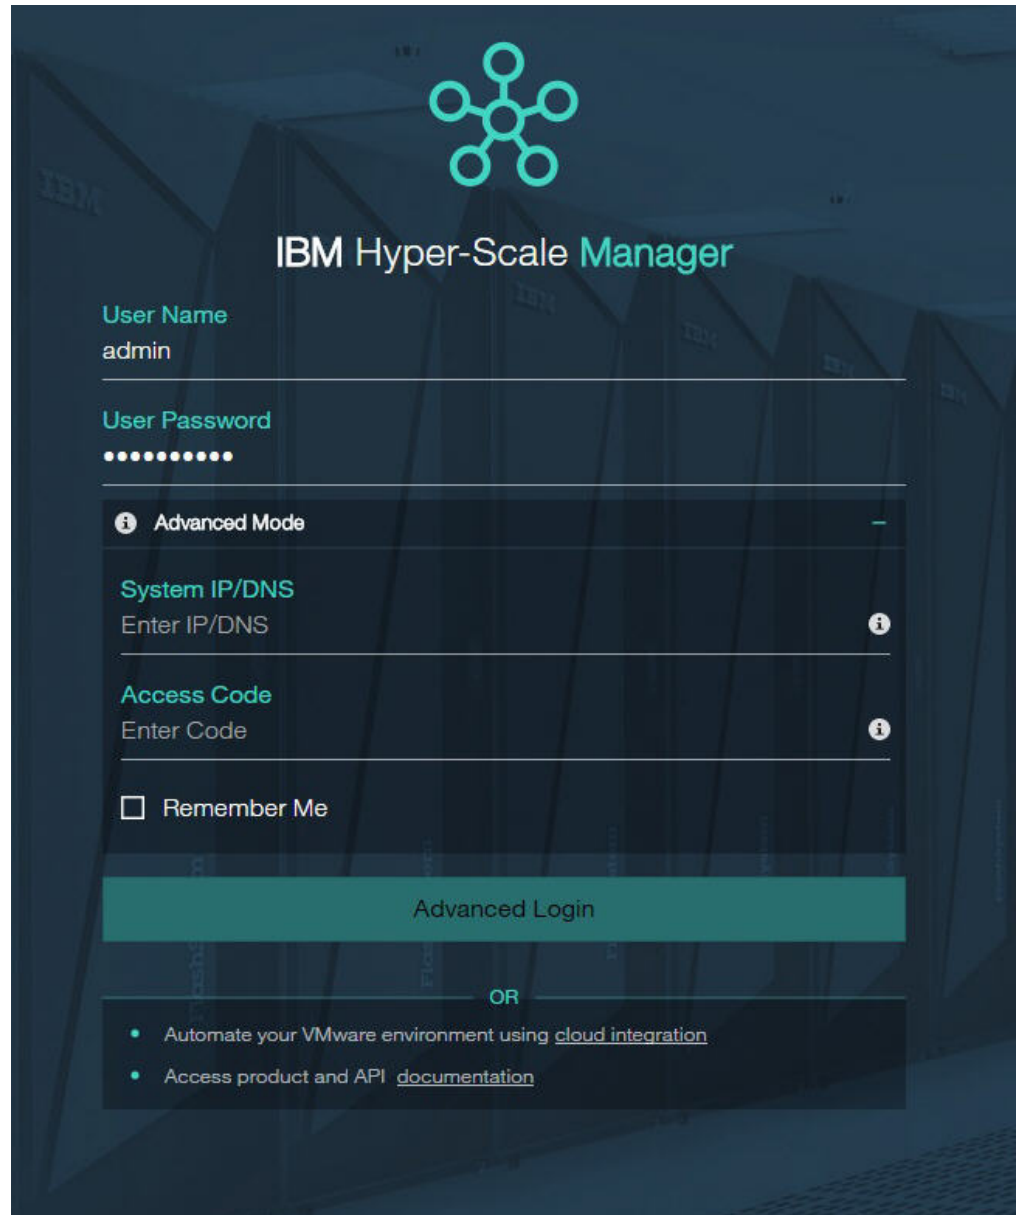


Figure 1. Advanced login screen

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:

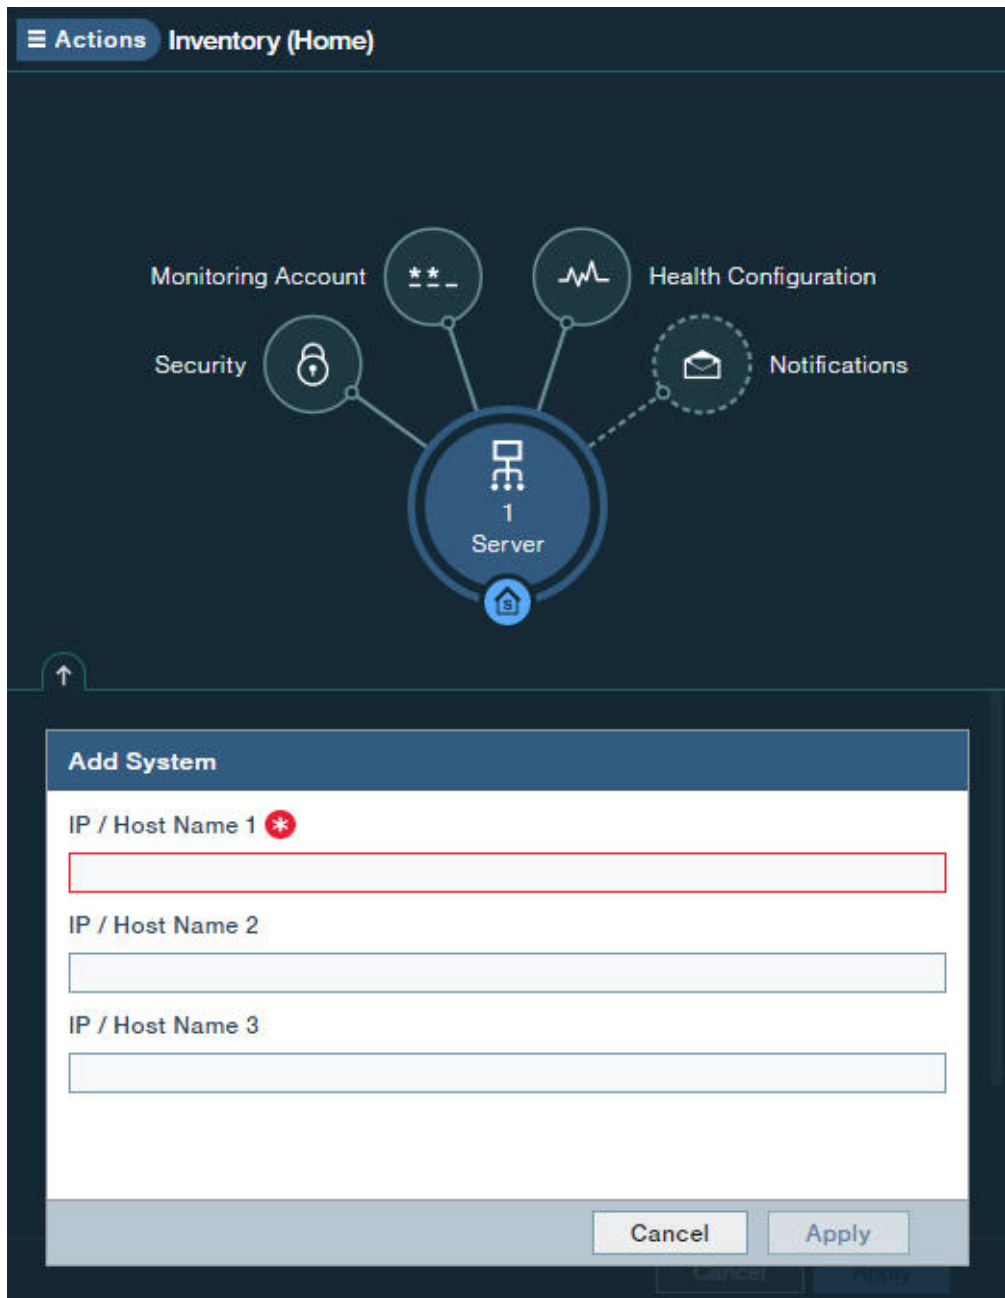


Figure 2. Adding a system to the inventory

3. Enter an IP/DNS address or Host Name for the system you want to add.
4. If you want to add an access code, click **Security** on the Hub to add the Management Server access code. The access code restricts administrative configuration actions to the Management Server when you log in.
5. Click **Monitoring Account** on the Hub to add or change the Monitoring Account information used by the Management Server to monitor the systems in your inventory.

Note: To apply the Monitoring Account to all of the systems in your inventory, select **Actions > Monitoring Account > Set Account on Systems**.

6. Click **Notifications** on the Hub to allow and add users to receive notifications about Hyper-Scale Manager issues to the destination email.
7. Click **Apply** to save the information entered. Click **Cancel** to exit without saving the information entered.

Replacing the IBM Hyper-Scale Manager certificate

This task 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. This task guides 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 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.x -----
-----
-----04/25/2016 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) Start/Stop Hyper-Scale Manager Service
10) 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.

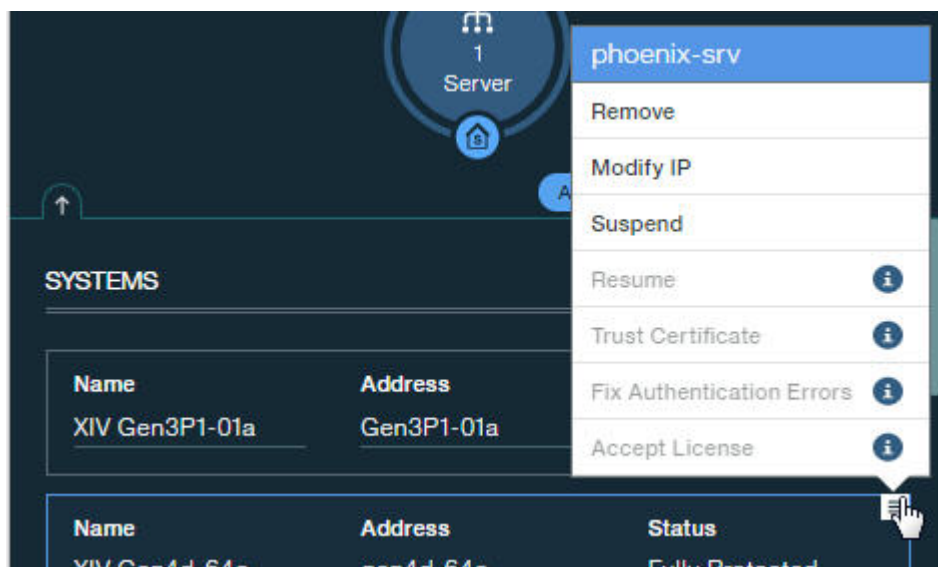


Figure 3. Viewing system certificate, licensing and authentication

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

Upgrading IBM Hyper-Scale Manager version 5.x

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

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 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 > 1  
  
IBM Hyper-Scale Manager will be upgraded from version 5.0.2 to <version>.  
Are you sure you want to upgrade now? [y/n] >
```

6. From the menu, enter the number that corresponds to Upgrade and Enter.
7. Type y to begin the upgrade.
8. 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.
```

9. Press Enter to view the license agreement. Once you accept the license, the following message is displayed.

In IBM's continuing efforts to improve our offerings, we would appreciate your allowing us access to your IBM Hyper-Scale Manager usage data. We intend to analyze this data to possibly enhance future offerings and to improve your client experience. If you agree, this information will be sent to IBM, will be used for the described purposes, and will be treated as IBM Confidential. Do you agree? [Y/n] >

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 IBM Confidential.

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.

10. The upgrade procedure begins. Continue working when the upgrade procedure completes successfully.

Migrating IBM Hyper-Scale Manager inventory from version 4.x

If you have previously worked with IBM Hyper-Scale Manager 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 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 "Setting up IBM Hyper-Scale Manager" 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.

Chapter 3. 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) Exit
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) Exit  
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
```

- 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 6 to configure your Monitoring Account information.

Chapter 4. Viewing the dashboard and system health score

The **Dashboard** displays an aggregated view of your system health, your data reduction rate and savings, provisioned capacity and performance metrics. This section focuses on your system health and its health score.

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

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 4. 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 affect 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. However, 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 below. Click on the Settings button in the Health widget to define your own health reduction settings, per severity, for how much is reduced from your total dashboard health score.

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 are not included in the health and do not affect the score.

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

- Two systems are connected, but not defined as *multipath*, thereby leaving the system at risk in case of a failure
- Host connectivity where the Host is connected to one module, but not set up for *multipathing*, thereby leaving the Host at risk in case of module failure
- Hosts and Volumes are not mapped and there is reclaimable capacity
- 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

Click on the alerts or tips to navigate to and fix the problematic object and improve your health score.

Chapter 5. Understanding and monitoring capacity

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. But, interpreting capacity in your storage environment can be surprisingly tricky. Especially when storage vendors describe their capacity in different ways. 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.

Managing and utilizing the storage system's physical and projected capacities is a key aspect of system management. The following figure illustrates capacity accounting, monitoring and planning.

On the right side of Figure 5, 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 the figure, the **Projected system capacity** is the combined calculation of the expected reduction you define with the expected accumulated data workload (in this case, x5.0). 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 5. Projected vs. Allocated vs. Physical

Figure 5 on page 19 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. First, 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. Next, make basic projections by using typical compression rates for calculating potential savings. **Over allocation** can also save even more space. See the *IBM XIV Gen3 Product Overview (GC27-3912)* and the *IBM FlashSystems A9000 or A9000R Product Overview (GC27-8583 or SC27-8558, respectively)* guides for more information on thin provisioning.

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 Projected Reduction Ratio keeps the Projected System Capacity calculation steady.

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 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 Dashboard in the IBM Hyper-Scale Manager UI.

Monitoring your capacity

About this task

From the IBM Hyper-Scale Manager UI, you can now monitor all of your systems' capacity from the dashboard across your "A" and "XIV" storage families at one time. 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.
2. Click on one of the tabs above the gauges to display the selected tier, as shown:

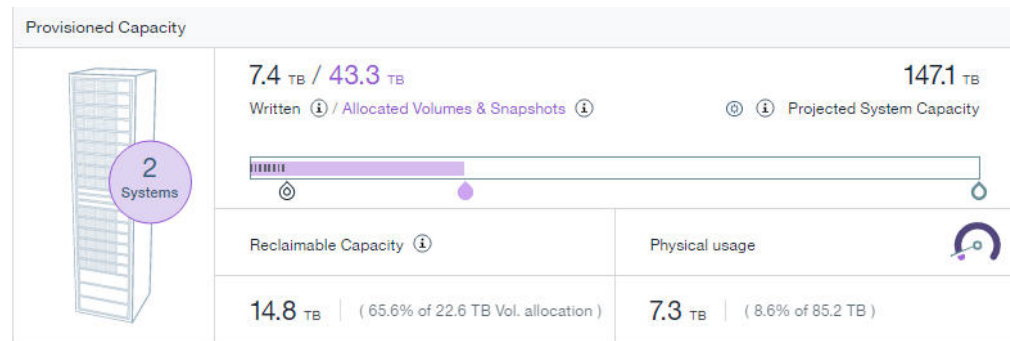


Figure 6. Select tier

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.

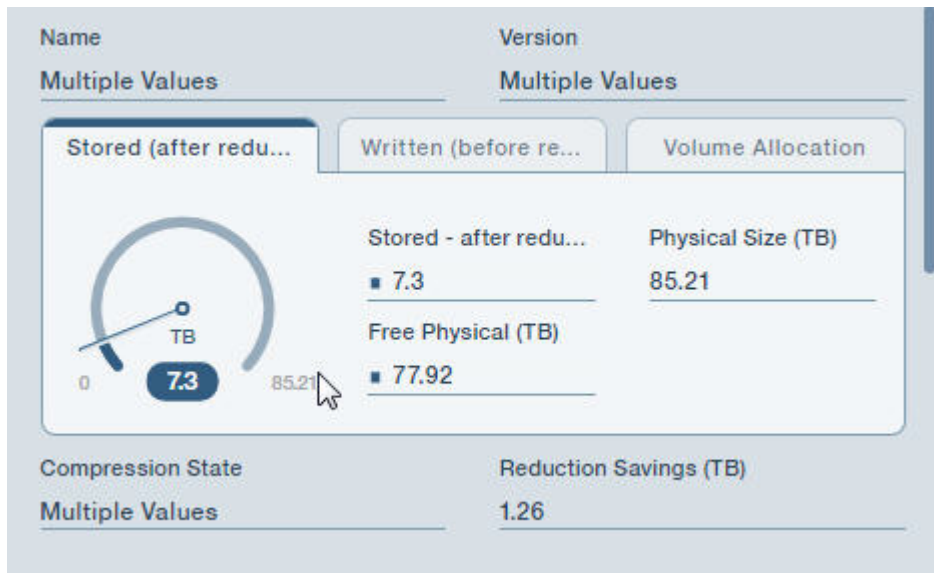


Figure 7. Capacity view

In IBM FlashSystems 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 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

You can generate a capacity planning report from 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 - System Usage (All System Types)
- Part 2 - IBM XIV Gen3/Spectrum Accelerate Domain Usage
- Part 3 - IBM XIV Gen3/Spectrum Accelerate Pool Usage

For all three parts (Systems, Domains and Pools), graphical representation illustrates the capacity usage over periods of time, usage growth rate in GB per week, and detailed graphs per object capacity usage, respectively.

For example, Figure 8 on page 25 shows an overall view for all of the systems in your inventory and a capacity forecast over the next 12 months

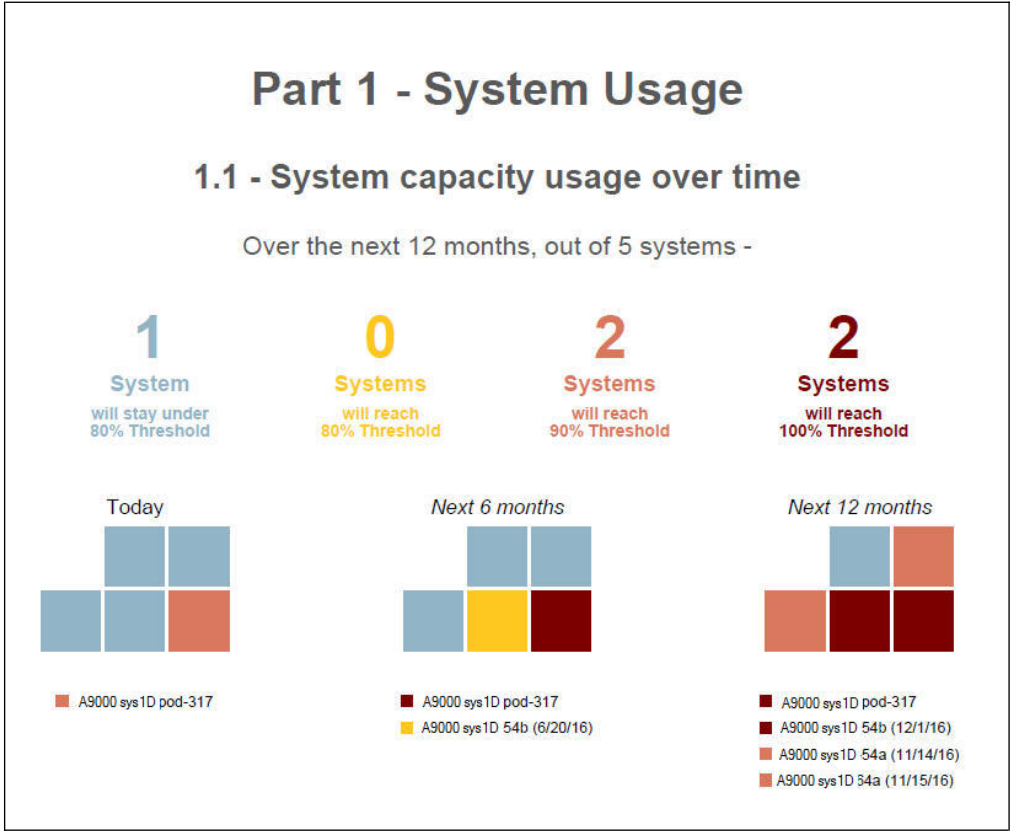


Figure 8. System capacity usage over time

The graph, Figure 9, ranks the system per used capacity growth rate.




Figure 9. System by usage growth rate

And for each system, individual graphs show the progression of the system capacity allocation growth. See Figure 10 on page 26.




Figure 10. 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.

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.

Moving capacity data between Management Server instances

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.

```

-----
----- IBM 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 (*.csv) 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 CSV 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 CSV 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.

```

-----
----- IBM 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 (*.csv) 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. The file is imported.

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.

```
-----  
----- IBM 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. 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, however, 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 **System Connectivity** form. If there are no Targets connected to this System, the display will be empty.
3. From the top of the **System Connectivity** form, click on the plus (+) sign to add a **Target**. The **Creating New Connectivity** form is displayed.

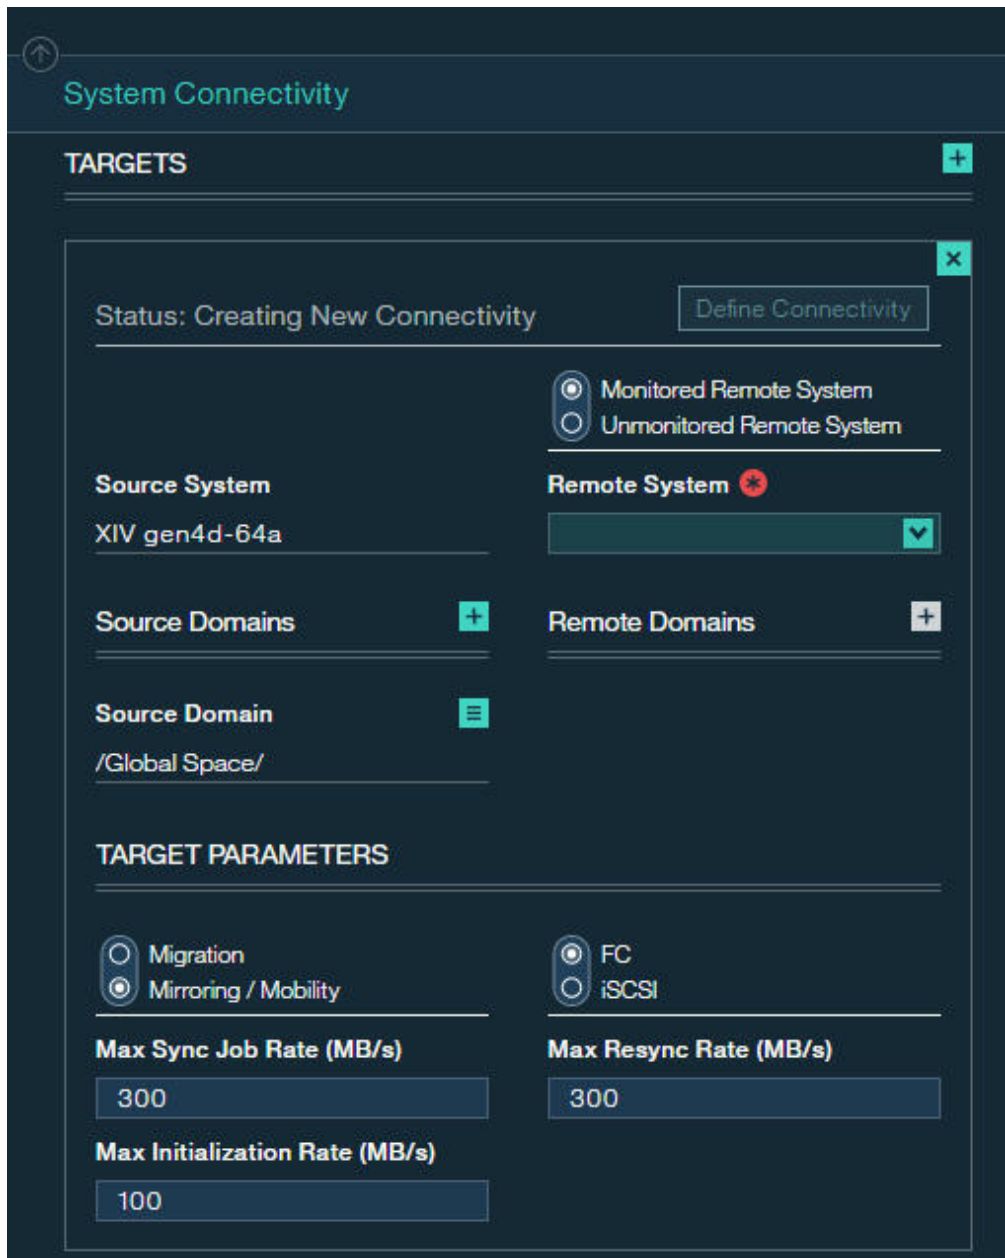


Figure 11. Creating new connectivity

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 **Connectivity Patch Panel** displaying 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.

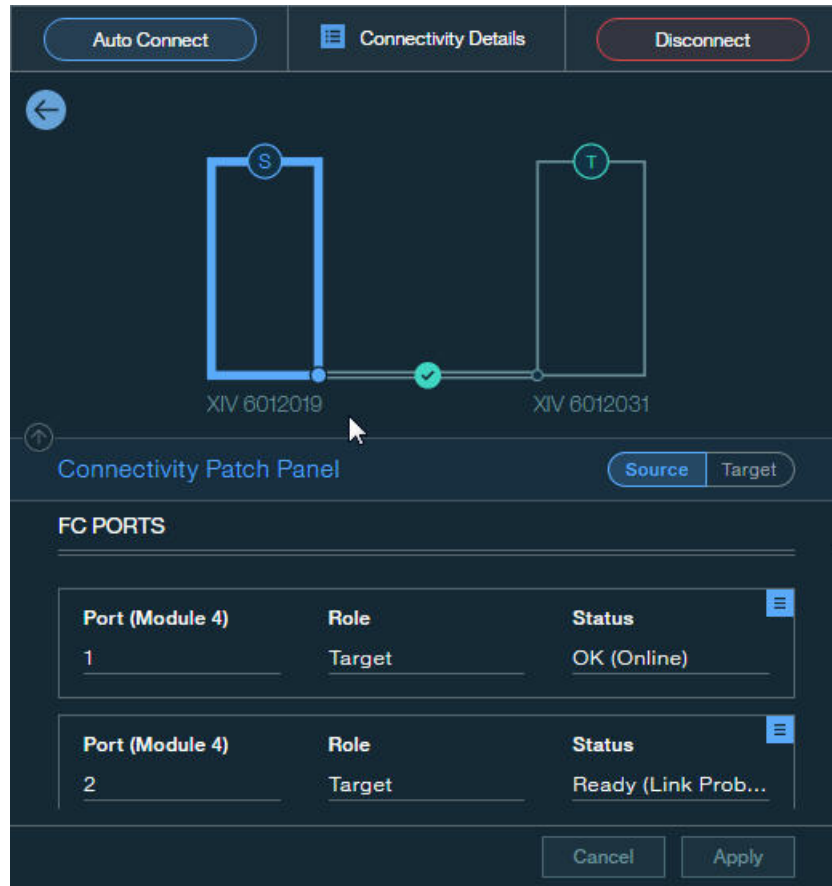


Figure 12. Connectivity patch panel

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 **Connectivity Patch Panel** to see the ports and the port status defined on both.
7. To see a graphical image of the **Connectivity Patch Panel** and the connections between the ports of the Source and Target systems, click on **Connectivity Details**.

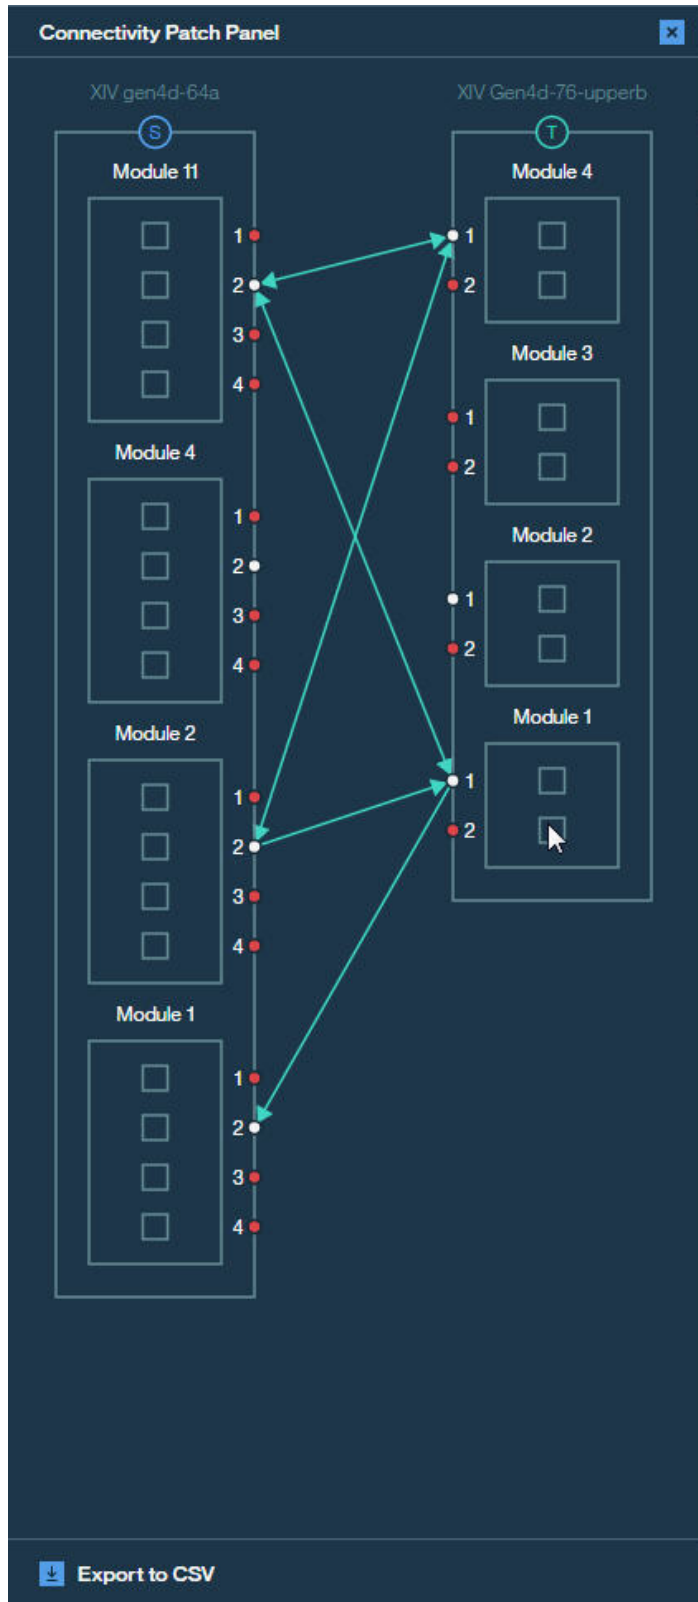


Figure 13. Connectivity details

Click **CSV** to download a more detailed report. A file can be found in the containing folder with a breakdown of 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.

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 Connectivity Patch Panel.

Procedure

1. Click **Add**. The **FC PORTS** or **ISCSI 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**.

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 24

A

Accounting

Capacity management 19

Understanding capacity 19

Adding Systems 6

analytics 28

B

Backing up IBM Hyper-Scale
Manager 13

backup 13

backup folder 28

C

capacity analytics 27, 28, 29

Capacity Management

Provisioning capacity

Accounting 19

Effective capacity 19

Physical capacity 19

Reclaimable capacity 19

Total allocation 19

Written by Hosts 19

capacity planning 24

Capacity Planning PDF Report 24

Capacity report

generating a report 23

certificate

of the IBM Hyper-Scale Manager 9,
10

Connectivity

Defining new connectivity 31

managing 31

D

Dashboard

health score 17

Domain usage 24

E

Exporting capacity data 27

F

forecast 24

forecasted capacity 24

G

Generating a capacity planning
report 24

H

how to back up 13

Hyper-Scale Manager

Setting up Hyper-Scale Manager 5

I

IBM Hyper-Scale Manager

Management server 3

Migrating your inventory from

4.x 12

User Interface (UI) 3

IBM Hyper-Scale Manager certificate 9,
10

Importing capacity data 28

Installation

IBM Hyper-Scale Manager 3

Management server 3

Introduction 1

inventory configuration and
preferences 14

L

Logging in

Advanced login 6

M

maintenance account 28

Management server

Installation 3

Management Server

Monitoring account 3

Migrating to 5.x 12

Monitoring account

Management Server 3

Monitoring and statistics 26

Monitoring capacity 21

P

PDF file 24

Pool usage 24

pools statistics 24

R

resetting the raw capacity data 29

restore 13

Restoring IBM Hyper-Scale Manager 14

S

score 17

SFTP 28

Standalone application 28

structure of the PDF file 24

System allocation 24

System capacity

planning 23

System connectivity 31

System Health 17

T

The report legend 24

threshold 24

U

Upgrading

IBM Hyper-Scale Manager 11

upload folder 28

V

Virtual appliance 28



Printed in USA

SC27-8560-04

