



Monitoring NVIDIA GPUs

"NVIDIA: GPU Monitoring" PowerPack version 102

Table of Contents

Introduction to Monitoring NVIDIA GPU Devices	3
What Does the NVIDIA: GPU Monitoring PowerPack Monitor?	4
Installing the NVIDIA: GPU Monitoring PowerPack	4
Configuring NVIDIA GPU Devices for Monitoring	6
Prerequisites for Monitoring NVIDIA GPU Devices	6
Creating an SSH/Key Credential for NVIDIA GPU Devices	7
Creating an SSH/Key Credential for NVIDIA GPU in the Skylar One Classic User Interface	8
Discovering NVIDIA GPU Devices	8
Aligning Dynamic Applications to NVIDIA GPU Devices	9
Configuring the NVIDIA GPU Device Template	9
Using the Device Template to Align Dynamic Applications to NVIDIA GPU Devices	10

Chapter

1

Introduction to Monitoring NVIDIA GPU Devices

Overview

This manual describes how to monitor NVIDIA graphics processing unit (GPU) devices in Skylar One using the "NVIDIA: GPU Monitoring" PowerPack.

This chapter covers the following topics:

<i>What Does the NVIDIA: GPU Monitoring PowerPack Monitor?</i>	4
<i>Installing the NVIDIA: GPU Monitoring PowerPack</i>	4

NOTE: ScienceLogic provides this documentation for the convenience of ScienceLogic customers. Some of the configuration information contained herein pertains to third-party vendor software that is subject to change without notice to ScienceLogic. ScienceLogic makes every attempt to maintain accurate technical information and cannot be held responsible for defects or changes in third-party vendor software. There is no written or implied guarantee that information contained herein will work for all third-party variants. See the End User License Agreement (EULA) for more information.

What Does the NVIDIA: GPU Monitoring PowerPack Monitor?

To monitor NVIDIA GPUs using Skylar One, you can install the "NVIDIA: GPU Monitoring" PowerPack. This PowerPack includes:

- Dynamic Applications that enable Skylar One to align to and monitor NVIDIA GPU devices
- Event policies that are triggered when NVIDIA GPU devices meet certain status criteria
- A sample SSH Credential ("NVIDIA GPU Monitoring - Example") that you can use to create your own NVIDIA credential
- A device template ("NVIDIA GPU Monitor Template") that aligns NVIDIA Dynamic Applications with devices

Installing the NVIDIA: GPU Monitoring PowerPack

Before completing the steps in this manual, you must import and install the latest version of the "NVIDIA: GPU Monitoring" PowerPack.

NOTE: In version 101 and later of the "NVIDIA: GPU Monitoring" PowerPack, GPU indexing uses the universally unique identifier (UUID) if a serial number is reported as "N/A" or unavailable. If you upgrade from version 100 of the PowerPack and your environment previously reported "N/A" serial numbers, you might notice changes in performance graphs and newly created GPU objects in configuration views. This behavior is expected, and data and graphs should fully normalize after two to three polling cycles.

TIP: By default, installing a new version of a PowerPack will overwrite all content in that PowerPack that has already been installed on the target system. You can use the **Enable Selective PowerPack Field Protection** setting in the **Behavior Settings** page (System > Settings > Behavior) to prevent the new version of the PowerPack from overwriting local changes for some commonly customized fields.

NOTE: For details on upgrading Skylar One, see the relevant [Skylar One Platform Release Notes](#).

To download and install the PowerPack:

1. Search for and download the PowerPack from the **PowerPacks** page at the [ScienceLogic Support Center](#) (Skylar One > PowerPacks, login required).
2. In Skylar One, go to the **PowerPacks** page (System > Manage > PowerPacks).
3. Click the actions icon and choose *Import PowerPack*. The **Import PowerPack** dialog box appears.

4. Click **[Browse]** and navigate to the PowerPack file from step 1.
5. Select the PowerPack file and click **[Import]**. The **PowerPack Installer** modal displays a list of the PowerPack contents.
6. Click **[Install]**. The PowerPack is added to the **PowerPacks** page.

NOTE: If you exit the **PowerPack Installer** modal without installing the imported PowerPack, the imported PowerPack will not appear in the **PowerPacks** page. However, the imported PowerPack will appear in the **Imported PowerPacks** modal. This page appears when you click the **[Actions]** menu and select *Install PowerPack*.

Chapter

2

Configuring NVIDIA GPU Devices for Monitoring

Overview

This section describes how to configure NVIDIA GPU devices for monitoring by Skylar One using the "NVIDIA: GPU Monitoring" PowerPack.

This chapter covers the following topics:

<i>Prerequisites for Monitoring NVIDIA GPU Devices</i>	6
<i>Creating an SSH/Key Credential for NVIDIA GPU Devices</i>	7
<i>Discovering NVIDIA GPU Devices</i>	8
<i>Aligning Dynamic Applications to NVIDIA GPU Devices</i>	9

Prerequisites for Monitoring NVIDIA GPU Devices

To configure Skylar One to monitor NVIDIA GPU devices using the "NVIDIA: GPU Monitoring" PowerPack, you must first have the following:

- SSH credentials with permissions to run the `nvidia-smi` command
- Physical NVIDIA GPU devices that you can align with the Dynamic Applications included in this PowerPack

Creating an SSH/Key Credential for NVIDIA GPU Devices

To configure Skylar One to monitor NVIDIA GPU devices, you must first create an SSH/Key credential, which allows the Dynamic Applications in the "NVIDIA: GPU Monitoring" PowerPack to communicate with NVIDIA GPU devices.

The PowerPack includes an example SSH/Key credential that you can edit for your own use.

To configure an SSH/Key credential to access NVIDIA GPU devices:

1. Go to the **Credentials** page (Manage > Credentials).
2. Locate the "NVIDIA GPU Monitoring - Example" sample credential, then click its actions icon (⋮) and select **Duplicate**. A copy of the credential appears, called "NVIDIA GPU Monitoring - Example copy."
3. Click the actions icon (⋮) for the "NVIDIA GPU Monitoring - Example copy" credential and select **Edit**. The **Edit Credential** modal page appears.
4. Supply values in the following fields:
 - **Name**. Enter a new name for the credential.
 - **Hostname/IP**. Type "%D". Skylar One replaces this variable with the device's IP address.
 - **Timeout (ms)**. Enter the time in milliseconds after which Skylar One stops trying to communicate with the monitored host.
 - **Username**. Enter the SSH account username that Skylar One uses to connect to the monitored host.
 - **Password**. Enter the password for the SSH account.
 - **Private Key (PEM Format)**. Enter the SSH private key.

NOTE: The private key can have a maximum of 64 characters per line. Therefore, you cannot use keys in the OpenSSH format, because that format uses 70 characters per line. When you attempt to save the credential, Skylar One validates that the private key entered is in the correct format. You can save the credential only if the private key is correctly formatted.

5. Click **[Save & Close]**.

Creating an SSH/Key Credential for NVIDIA GPU in the Skylar One Classic User Interface

To configure an SSH/Key credential to access an NVIDIA GPU device using the classic Skylar One user interface:

1. Go to the **Credential Management** page (System > Manage > Credentials).
2. Locate the "NVIDIA GPU Monitoring - Example" credential, then click its wrench icon (). The **Edit SSH/Key Credential** modal appears.
3. Complete the following fields:
 - **Name**. Enter a new name for the credential.
 - **Hostname/IP**. Type "%D". Skylar One replaces this variable with the device's IP address.
 - **Timeout (ms)**. Enter the time in milliseconds after which Skylar One stops trying to communicate with the monitored host.
 - **Username**. Enter the SSH account username that Skylar One uses to connect to the monitored host.
 - **Password**. Enter the password for the SSH account.
 - **Private Key (PEM Format)**. Enter the SSH private key.

NOTE: The private key can have a maximum of 64 characters per line. Therefore, you cannot use keys in the OpenSSH format, because that format uses 70 characters per line. When you attempt to save the credential, Skylar One validates that the private key entered is in the correct format. You can save the credential only if the private key is correctly formatted.

4. Click the **[Save As]** button.

Discovering NVIDIA GPU Devices

To create and run a discovery session that discovers an NVIDIA GPU root device, perform the following steps:

1. On the **Devices** page () or the **Discovery Sessions** page (Devices > Discovery Sessions), click the **[Add Devices]** button. The **Select** page appears.
2. Click the **[Unguided Network Discovery]** button. Additional information about the requirements for discovery appears in the **General Information** pane to the right.
3. Click **[Select]**. The discovery wizard appears, starting with the **Basic Information** page.
4. Complete the following fields:
 - **Name**. Enter a unique name for this discovery session. This name displays in the list of discovery sessions on the **[Discovery Sessions]** tab.

- **Description.** Optional. Enter a short description of the discovery session. You can use the text in this description to search for the discovery session on the **[Discovery Sessions]** tab.
 - **Select the organization to add discovered devices to.** Select the name of the organization to which you want to add the discovered devices.
5. Click **[Next]**. The **Credential Selection** page of the **Add Devices** wizard appears.
 6. On the **Credential Selection** page, locate and select the [SSH/ Key credential](#) you created.
 7. Click **[Next]**. The **Discovery Session Details** page of the **Add Devices** wizard appears.
 8. Complete the following fields:
 - **List of IPs/Hostnames.** Enter the IP address for the NVIDIA GPU root device.
 - **Which collector will discover these devices?** Select an existing collector to monitor the discovered devices.
 - **Run after save.** Select this option if you want to run this discovery session as soon as you click **[Save and Close]**.
 9. Click **[Save and Close]** to save the discovery session. The **Discovery Sessions** page displays the new discovery session.
 10. If you selected the **Run after save** option, the discovery session runs immediately, and the **Discovery Logs** page displays any relevant log messages. If the discovery session locates and adds any devices, the **Discovery Logs** page includes a link to the **Device Investigator** page for the discovered device.

Aligning Dynamic Applications to NVIDIA GPU Devices

You can use a **device template** to save a device configuration and apply it to multiple devices. The "NVIDIA: GPU Monitoring" PowerPack includes the "NVIDIA GPU Monitor Template," which enables Skylar One to align all Dynamic Applications to the root component device.

Configuring the NVIDIA GPU Device Template

Before you can use the "NVIDIA GPU Monitor Template," you need to configure the template so that each Dynamic Application in the template aligns with the [credential you created earlier](#).

To configure the device template:

1. Go to the **Configuration Templates** page (Devices > Templates).
2. Locate the "NVIDIA GPU Monitor Template" and click its wrench icon (). The **Device Template Editor** modal appears.
3. Change the name of the template and click **[Save As]**. This creates a copy of the template.
4. Click the **[Dyn Apps]** tab. The **Editing Dynamic Application Subtemplates** page appears.
5. In the **Credentials** drop-down field, select the credential that you created for NVIDIA GPU devices.
6. Click the next Dynamic Application listed in the **Subtemplate Selection** section on the left side of the page and then select the credential you created in the **Credentials** field.

7. Repeat step 6 until you have selected the appropriate credential for all of the Dynamic Applications listed in the **Subtemplate Selection** section.
8. Click **[Save]**.

Using the Device Template to Align Dynamic Applications to NVIDIA GPU Devices

To align the NVIDIA GPU Dynamic Applications to NVIDIA GPU devices:

1. Go to the **Device Manager** page (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface).
2. Select the checkbox for all devices you want to align with the NVIDIA GPU Dynamic Applications.
3. In the **Select Actions** field, select the option *MODIFY by Template* and then click the **[Go]** button. The **Device Template Editor** modal appears.
4. Complete the following fields:
 - In the **Template** drop-down field, select the name of *the device template you configured earlier*.
 - In the **Credentials** drop-down field, select *the credential you created earlier*.
5. Click the **[Apply]** button, and then click **[Confirm]** to align the Dynamic Applications to the selected devices.
6. Confirm that the Dynamic Applications were aligned with the selected devices by clicking on a device's wrench icon (🔧) and then clicking the **[Collections]** tab. The page lists any Dynamic Applications aligned with the devices.

© 2003 - 2026, ScienceLogic, Inc.

All rights reserved.

ScienceLogic™, the ScienceLogic logo, and ScienceLogic's product and service names are trademarks or service marks of ScienceLogic, Inc. and its affiliates. Use of ScienceLogic's trademarks or service marks without permission is prohibited.

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information herein, the information provided in this document may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described herein at any time without notice.



800-SCI-LOGIC (1-800-724-5644)

International: +1-703-354-1010