ScienceLogic

Monitoring NVIDIA GPU

NVIDIA GPU PowerPack version 100

Table of Contents

Introduction	3
What Does the NVIDIA GPU PowerPack Monitor?	3
Installing the NVIDIA GPU PowerPack	4
Configuration	5
Prerequisites for Monitoring NVIDIA GPU Devices	5
Creating an SSH/Key Credential for NVIDIA GPU	5
Creating an SSH/Key Credential for NVIDIA GPU in the SL1 Classic User Interface	6
Aligning Dynamic Applications to NVIDIA GPU Devices	7
Configuring the Device Template	7
Using the Device Template to Align Dynamic Applications to NVIDIA GPU Devices	8

Chapter

Introduction

Overview

This manual describes how to monitor NVIDIA GPU devices in SL1 using the Dynamic Applications in the NVIDIA GPU PowerPack.

The following sections provide an overview of NVIDIA GPU and the NVIDIA GPU PowerPack:

This chapter covers the following topics:

What Does the NVIDIA GPU PowerPack Monitor?	. 3
Installing the NVIDIA GPU PowerPack	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 PowerPack Monitor?

To monitor NVIDIA GPUs using SL1, you can install the NVIDIA GPU PowerPack. This PowerPack enables you to align Dynamic Applications to NVIDIA GPU devices to collect data.

The NVIDIA GPU PowerPack includes:

• Dynamic Applications that enable SL1 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 PowerPack

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

TIP: By default, installing a new version of a PowerPack overwrites all content from a previous version of 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 new PowerPacks from overwriting local changes for some commonly customized fields. For more information, see the section on *Global Settings*.

To download and install the PowerPack:

- Search for and download the PowerPack from the PowerPacks page (Product Downloads > PowerPacks & SyncPacks) at the <u>ScienceLogic Support Site</u>.
- 2. In SL1, go to the **PowerPacks** page (System > Manage > PowerPacks).
- 3. Click the [Actions] button 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

Configuration

Overview

The following sections describe how to configure and discover NVIDIA GPU devices for monitoring by SL1 using the NVIDIA GPU PowerPack:

This chapter covers the following topics:

Prerequisites for Monitoring NVIDIA GPU Devices	5
Creating an SSH/Key Credential for NVIDIA GPU	5
Aligning Dynamic Applications to NVIDIA GPU Devices	7

Prerequisites for Monitoring NVIDIA GPU Devices

To configure the SL1 system to monitor NVIDIA GPU devices using the NVIDIA GPU PowerPack, you must first have the following information about NVIDIA GPU:

- SSH (Secure Shell) 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

To configure SL1 to monitor NVIDIA GPU devices, you must first create an SSH/Key credential. This credential allows the Dynamic Applications in the *NVIDIA GPU* 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 an NVIDIA GPU:

- 1. Go to the **Credentials** page (Manage > Credentials).
- 2. Locate the Nvidia GPU Monitoring- Example sample credential, click its [Actions] icon (--) and select *Duplicate*. A copy of the credential, called Nvidia GPU Monitoring Example copy appears.
- 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. Type a new name for the credential.
 - Hostname/IP. Type "%D". SL1 will replace it with the device's IP.
 - *Timeout (ms)*. The time in milliseconds, after which SL1 will stop trying to communicate with the monitored host.
 - Username. Type the SSH account username. This will be used to connect to the monitored host.
 - Password. Type the password for the SSH account.
 - Private Key (PEM Format). Type 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, SL1 will validate that the private key entered is in the correct format. You will be able to 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 SL1 Classic User Interface

To configure SL1 to monitor NVIDIA GPU devices, you must first create an SSH/Key credential. This credential allows the Dynamic Applications in the *NVIDIA GPU* PowerPack to communicate with NVIDIA GPU devices.

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

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

- 1. Go to the Credential Management page (System > Manage > Credentials).
- 2. Locate the Nvidia GPU Monitoring- Example credential, then click its wrench icon (*P*). The Edit SSH/Key Credential modal page appears:

- 3. Complete the following fields:
 - Name. Type a new name for the credential.
 - Hostname/IP. Type "%D". SL1 will replace it with the device's IP.
 - **Timeout (ms)**. The time in milliseconds, after which SL1 will stop trying to communicate with the monitored host.
 - Username. Type the SSH account username. This will be used to connect to the monitored host.
 - **Password**. Type the password for the SSH account.
 - Private Key (PEM Format). Type 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, SL1 will validate that the private key entered is in the correct format. You will be able to save the credential only if the private key is correctly formatted.

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

Aligning Dynamic Applications to NVIDIA GPU Devices

A *device template* allows you to save a device configuration and apply it to multiple devices. The *NVIDIA GPUPowerPack* includes the "Nvidia GPU Monitor Template" which enables SL1 to align all Dynamic Applications to the root component device.

Configuring the 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 (*P*). The **Device Template Editor** modal page appears.
- 3. Change the name of the template and click [Save As]. This will create a copy of the template.
- 4. Click the [Dyn Apps] tab. The Editing Dynamic Application Subtemplates page appears.
- 5. In the Credentials drop-down list, select the credential that you created for NVIDIA GPU.
- 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 5 until you have selected that credential in the **Credentials** field 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 > Device Manager).
- 2. On the **Device Manager** page, select the checkbox for all devices where you want to align the NVIDIA GPU Dynamic Applications.
- 3. In the **Select Actions** field, in the lower right, select the option MODIFY by Template and click the **[Go]** button. The **Device Template Editor** page appears:
- 4. Complete the following fields:
 - In the *Template* drop-down list, select the name of the device template you configured earlier.
 - In the Credentials drop-down list, 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 (*P*) and selecting the **[Collections]** tab. Any aligned Dynamic Applications will be listed.

© 2003 - 2024, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

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 on this Site, information on this Site 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 in this Site at any time without notice.

Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and [™] symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic[™]
- EM7[™] and em7[™]
- Simplify IT[™]
- Dynamic Application™
- Relational Infrastructure Management[™]

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: legal@sciencelogic.com. For more information, see https://sciencelogic.com.



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

International: +1-703-354-1010