ScienceLogic

PagerDuty SyncPack

Version 1.0.2

Table of Contents

Chapter

1

Introduction to the PagerDuty SyncPack

Overview

This chapter describes how you can use the "PagerDuty" SyncPack to automatically trigger, acknowledge, and resolve SL1 events and PagerDuty incidents between PagerDuty and SL1 . The integration is bi-directional, from SL1 to PagerDuty.

This SyncPack uses the "PagerDuty Incident Integration" PowerPack.

This chapter covers the following topics:

What Can I Do with this SyncPack?

The "PagerDuty" SyncPack lets you automatically trigger, acknowledge, and resolve SL1 events and PagerDuty incidents between PagerDuty and SL1. The integration is bidirectional, from SL1 to PagerDuty.

You can configure this SyncPack to perform the following integrations:

- Create an alert in PagerDuty from an SL1 event.

- Collect on-call details from PagerDuty and add it to the notes of the corresponding SL1 event.

- Update a synced Opsgenie alert when an event is acknowledged in SL1, including closing an alert in Opsgenie that was resolved by a corresponding SL1 event.

- Collect resolved alert data from Opsgenie and clear the corresponding SL1 event.

Contents of the SyncPack

This section lists the contents of the "PagerDuty" SyncPack.

PowerFlow Applications

- *Create Event in PagerDuty from SL1 Event*. This application sends SL1 event information to PagerDuty. The application translates and formats the event data from SL1 so that it can be consumed by PagerDuty to trigger, acknowledge, or resolve an incident.

- *Get PagerDuty log and incident data*. This application queries the PagerDuty API to collect log entries and incident data. The collected entries are formatted and passed to subordinate applications that process the data to update, acknowledge, and clear events in SL1.

- *Clear resolve event from ScienceLogic*. This application clears an event in SL1 that was resolved in PagerDuty. This application is called by the "Create resolve alert data from PagerDuty" application for each event that needs to be cleared.

- *Post event updates to ScienceLogic*. This application updates SL1 events using the data collected from the PagerDuty API.

For more information about how to configure these applications, see *Configuring and Aligning the PagerDuty Applications*.

PowerFlow Applications (Internal)

The following applications are "internal" applications that should not be run directly, but are automatically run by applications from the previous list:

- *Create resolve alert data from PagerDuty*. This application is called by the "Get PagerDuty log and incident data" application to process log messages that indicate that an incident has been resolved. This application queries the PagerDuty API for the alerts associated with the resolved incidents. If the alert was triggered from an SL1 event, the application calls the "Clear resolve event from ScienceLogic" application to clear the event.

- *Get trigger alert data from PagerDuty*. This application is called by the "Get PagerDuty log and incident data" application to process log messages that indicate that an incident has been triggered. This application queries the PagerDuty API for the alerts associated with those triggered incidents. If the alert was triggered from an SL1 event, the application calls the "Post event updates to ScienceLogic" application to add to incident ID to the event.

- *Get alert and user data from PagerDuty and ScienceLogic*. This application is called by the "Get PagerDuty log and incident data" application to process log messages that indicate that an incident has been acknowledged. This application queries the PagerDuty API for the alerts associated with acknowledged incidents. If the alert was triggered from an SL1 event, the application updates the SL1 event so that the acknowledgment matches PagerDuty. This application also determines the date of the most recently processed alert and stores that value to be used by the next execution of the "Get PagerDuty log and incident data" application.

Configuration Object

- *PagerDuty Base Config*. This configuration object can be used as a template after the SyncPack is installed on the PowerFlow system. The configuration object includes the following:

  - Details for connecting to the SL1 API, including the URL, username, and password

  - Details for connecting to the PagerDuty API, including the URL, username, and password

  - Details for connecting to the PagerDuty REST API, including the URL and API key

  - Details for connecting to an HTTP Proxy when communicating with PagerDuty, including the hostname or IP address, username, and password

  - An optional identifier that appends the event ID to create a PagerDuty dedup key. This can be used to avoid routing collisions when multiple SL1 platforms connect to the same PagerDuty instance.

  - Mapping between SL1 events and PagerDuty incidents including, event summary formats, event data, and indications to resolve or link incidents between SL1 and PagerDuty

Steps

The following steps are included in this SyncPack:

- Acknowledge the alert by user
- Check for since parameter after query
- Check conditional
- Get alert count and acknowledge the alert by user
- Get incident and user data
- Get resolve incident data
- Get trigger incident data
- Get user email from user details
- Convert ScienceLogic event data to PagerDuty event

Chapter

2

Installing the PagerDuty SyncPack

Overview

This chapter describes how to install the "PagerDuty" SyncPack, and how to install the "PagerDuty Incident Integration" PowerPack.

This chapter covers the following topics:

Prerequisites for this SyncPack

This SyncPack requires the following:

- The latest "PagerDuty Incident Integration" PowerPack
- Administrator access to PagerDuty

For the latest System Requirements, see the *SL1 PowerFlow SyncPack Release Notes*.

The following table lists the port access required by PowerFlow and this SyncPack:

Source IP	PowerFlow Destination	PowerFlow Source Port	Destination Port	Requirement
PowerFlow	SL1 API	Any	TCP 443	SL1 API Access
PowerFlow	PagerDuty API	Any	TCP 443	PagerDuty API Access
PowerFlow	SL1 Database	Any	TCP 7706	SL1 Database Access

Downloading and Installing the PagerDuty SyncPack

A SyncPack file has the **.whl** file extension type. You can download the SyncPack file from the ScienceLogic Support site.

Downloading the SyncPack

> **NOTE:** If you are installing or upgrading to the latest version of this SyncPack in an offline deployment, see *Installing or Upgrading in an Offline Environment* to ensure you install any external dependencies.

To locate and download the SyncPack:

1. Go to the ScienceLogic Support Site at https://support.sciencelogic.com/s/.
2. Click the **[Product Downloads]** tab and select *PowerPacks & SyncPacks*.
3. In the **Search** field, search for the SyncPack and select it from the search results. The **Release Version** page appears.
4. On the **[Files]** tab, click the down arrow next to the SyncPack version that you want to install, and select *Show File Details*. The **Release File Details** page appears.
5. Click the **[Download File]** button to download the SyncPack.

After you download the SyncPack, you can import it to your PowerFlow system using the PowerFlow user interface.

Importing the SyncPack

To import a SyncPack in the PowerFlow user interface:

1. On the **SyncPacks** page (⊚) of the PowerFlow user interface, click **[Import SyncPack]**. The **Import SyncPack** page appears.

2. Click **[Browse]** and select the **.whl** file for the SyncPack you want to install. You can also drag and drop a **.whl** file to the **Import SyncPack** page.

3. Click **[Import]**. PowerFlow registers and uploads the SyncPack. The SyncPack is added to the **SyncPacks** page.

4. You will need to activate and install the SyncPack in PowerFlow. For more information, see the following topic.

> **NOTE:** You cannot edit the content package in a SyncPack published by ScienceLogic. You must make a copy of a ScienceLogic SyncPack and save your changes to the new SyncPack to prevent overwriting any information in the original SyncPack when upgrading.

Installing the SyncPack

To activate and install a SyncPack in the PowerFlow user interface:

1. On the **SyncPacks** page of the PowerFlow user interface, click the **[Actions]** button (⋮) for the SyncPack you want to install and select *Activate & Install*. The **Activate & Install SyncPack** modal appears.

> **NOTE:** If you try to activate and install a SyncPack that is already activated and installed, you can choose to "force" installation across all the nodes in the PowerFlow system.

> **TIP:** If you do not see the PowerPack that you want to install, click the Filter icon (▽) on the **SyncPacks** page and select *Toggle Inactive SyncPacks* to see a list of the imported PowerPacks.

2. Click **[Yes]** to confirm the activation and installation. When the SyncPack is activated, the **SyncPacks** page displays a green check mark icon (✓) for that SyncPack. If the activation or installation failed, then a red exclamation mark icon (❗) appears.

3. For more information about the activation and installation process, click the check mark icon (✓) or the exclamation mark icon (❗) in the **Activated** column for that SyncPack. For a successful installation, the "Activate & Install SyncPack" application appears, and you can view the Step Log for the steps. For a failed installation, the **Error Logs** window appears.

4. If you have other versions of the same SyncPack on your PowerFlow system, you can click the **[Actions]** button (⋮) for that SyncPack and select *Change active version* to activate a different version other than the version that is currently running.

Installing the PagerDuty Incident Integration PowerPack

The "PagerDuty Incident Integration" PowerPack includes automation policies that align incidents triggered and updated in PagerDuty to events in SL1. The PowerPack also includes an automation action that is executed on the SL1 Database Server.

> **NOTE:** The "PagerDuty Incident Integration" PowerPack requires SL1 version 10.1.0 or later. For details on upgrading SL1, see the *SL1 Platform Release Notes*.

> **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 *System Administration* manual.)

To download and install the PowerPack:

1. Search for and download the PowerPack from the **PowerPacks** page (Product Downloads > PowerPacks & SyncPacks) at the ScienceLogic Support Site.

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

3

Configuring PowerFlow Applications for the PagerDuty SyncPack

Overview

This chapter describes how to set up the run book automations in SL1 and how to configure the PowerFlow applications in the "PagerDuty" SyncPack.

This chapter covers the following topics:

Workflow for Configuring the SyncPack

The following workflows describe how to configure SL1 and PowerFlow to work with the "PagerDuty" SyncPack.

Configuring SL1

1. *Create a Basic/Snippet credential to access PowerFlow*
2. *Edit the PagerDuty run book actions*
3. *Enable the PagerDuty run book automations*

Configuring PowerFlow

1. *Create a configuration object*
2. *Align the configuration object and configure the PagerDuty PowerFlow applications*
3. *Schedule the PowerFlow applications*

Configuring SL1

The following topics cover how to set up your SL1 instance to work with the "PagerDuty" SyncPack:

1. *Create a Basic/Snippet credential to access PowerFlow*
2. *Edit the PagerDuty run book actions*
3. *Enable the PagerDuty run book automations*

Creating a Basic/Snippet Credential to Access SL1 PowerFlow

You will need to create a Basic/Snippet credential in SL1 so that the action policies included in the "PagerDuty Incident Integration" PowerPack can access your PowerFlow system.

To create a Basic/Snippet credential:

1. In SL1, go to the **Credentials** page (Manage > Credentials).
2. Click **[Create New]** and select *Create Basic/Snippet Credential*. The **Create Credential** window appears.
3. Complete the following fields:

    - *Name*. Name of the credential. Can be any combination of alphanumeric characters. This field is required.
    - *All Organizations*. Toggle on (blue) to align the credential to all organizations, or toggle off (gray) and then select one or more specific organizations from the *What organization manages this service?* drop-down field to align the credential with those specific organizations.
    - *Timeout (ms)*. Time, in milliseconds, after which SL1 will stop trying to communicate with the authenticating server.

- *Hostname/IP*. Hostname or IP address of your PowerFlow system. This field is required.

- *Port*. Type "443". This field is required.

- *URL*. Type the URL for your PowerFlow system.

- *Username*. Username for a user account on the PowerFlow system.

- *Password*. Password for a user account on the PowerFlow system.

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

5. Take note of the SL1-assigned ID number for the new credential on the **Credentials** page, in the **ID** column. You will need the ID number when editing the input parameters of the run book actions included in the PowerPack, below.

NOTE: When you define a Basic/Snippet credential, the credential will automatically be aligned with the organization(s) where you are a member. To learn more about credentials and organizations, see the *Discovery and Credentials* manual

---

Editing the PagerDuty Run Book Action

The "PagerDuty Incident Integration" PowerPack includes the "PagerDuty: Trigger/Acknowledge/Resolve Alert" run book action. This action uses the "Run Integration Service Application" action type to trigger thePowerFlow application that sends events to PagerDuty. You can specify the credential ID and an optional platform identifier in a JSON structure that you enter in the *Input Parameters* field in the **Action Policy Editor** modal.

To edit the "PagerDuty: Trigger/Acknowledge/Resolve Alert" run book action:

1. In SL1, go to the **Actions** page (Registry > Run Book > Actions).

2. Locate the "PagerDuty: Trigger/Acknowledge/Resolve/Alert" automation action, and then click its wrench icon ( ). The **Editing Action** page appears.

3. In the *Input Parameters* field, change the values of the following parameters:

    - *credential_id*. Change the value to the credential ID that you noted earlier when creating a credential for your PowerFlow system in the previous procedure. This parameter is required.

    - *include_event*. Leave the value as "true".

    - *application_name*. Leave the value as "send_sl1_events_to_pagerduty".

    - *platform_identifier*. Change the value to the unique identifier you created for each event sent to PagerDuty. This value must match the "platform_identifier" parameter in the PowerFlow configuration object that you configured for this SL1 platform.

NOTE: For more information on configuring the *platform_identifier* field in your PowerFlow applications, see *Scheduling the PowerFlow Applications.*

---

4. Make sure the **Action State** is set to *Enabled*, and then click **[Save]**.

Enabling the PagerDuty Run Book Automations

The "PagerDuty Incident Integration" PowerPack includes the following automation policies:

- "PagerDuty": Acknowledge Alert
- "PagerDuty": Resolve Alert
- "PagerDuty": Trigger Alert

These policies update the PagerDuty incident with the state of the event in SL1. When the event is first detected in SL1, the "triggered" state is sent. When the event is acknowledged in SL1, the "acknowledged" state is sent. When the event is cleared in SL1, the "resolved" state is sent.

To enable the run book automations:

1. In SL1, go to the **Automation** page (Registry > Run Book > Automation).
2. Locate a PagerDuty automation policy and click its wrench icon (🔧). The **Automation Policy Editor** page appears.
3. Update the following fields:
   - **Policy State**. Select *Enabled*.
   - **Policy Priority**. Select *High* to ensure that this PowerFlow automation policy is added to the top of the queue.
   - **Available Actions**. If it is not already selected, select the "Run Integration Service Application" action and click the arrows to move it to **Aligned Actions**.

> **WARNING:** ScienceLogic highly recommends that you do not make changes to the **Policy Type**, **Repeat Time**, or **Align With** fields or the *And event is NOT acknowledged* setting.

4. Click **[Save]**.
5. Repeat steps 2-4 for the remaining PagerDuty run book automation policies.

Configuring PowerFlow

The following topics cover how to set up your PowerFlow instance to work with the "PagerDuty" SyncPack:

1. *Create a configuration object*
2. *Align the configuration object and configure the PagerDuty PowerFlow applications*
3. *Schedule the PowerFlow applications*

Creating a Configuration Object

A *configuration object* supplies the login credentials and other required information needed to execute the steps for a PowerFlow application. The **Configurations** page (⚙) of the PowerFlow user interface lists all available configuration objects for that system.

You can create as many configuration objects as you need. A PowerFlow application can only use one configuration object at a time, but you can use (or "align") the same configuration object with multiple applications.

For this SyncPack, you can make a copy of the "PagerDuty Base Config" configuration object, which is the sample configuration file that was installed with the "PagerDuty" SyncPack. Then you can integrate one PagerDuty instance with multiple SL1 instances by creating a separate configuration for each SL1 instance.

> **TIP:** The "PagerDuty Base Config" configuration object contains all of the required variables. Simply update the variables from that object to match your SL1 and PagerDuty settings.

To create a configuration object based on the "PagerDuty Base Config" configuration object:

1. In the PowerFlow user interface, go to the **Configurations** page (⚙).
2. Click the **[Actions]** button (⋮) for the "PagerDuty Base Config" configuration object and select *Edit*. The **Configuration** pane appears.
3. Click **[Copy as]**. The **Create Configuration** pane appears.
4. Complete the following fields:

   - *Friendly Name*. Name of the configuration object that will display on the **Configurations** page.
   - *Description*. A brief description of the configuration object.
   - *Author*. User or organization that created the configuration object.
   - *Version*. Version of the configuration object.

5. In the **Configuration Data** field, update the default variable definitions to match your PowerFlow configuration:

   - *event_api*. The URL of the PagerDuty event API for all PagerDuty instances. The default value is "https://events.pagerduty.com/v2/enqueue". You do not need to update this field.
   - *integration_key*. The routing/integration key to use when sending events.

     > **NOTE:** For information on creating a PagerDuty service and accessing the integration key, see https://support.pagerduty.com/docs/services-and-integrations#create-a-new-service.

   - *pd_base_url*. The base URL of the PagerDuty REST API for all PagerDuty instances. The default value is "https://api.pagerduty.com". You do not need to update this field.

- **rest_api_key**. The api key to use when syncing information to SL1 from PagerDuty.

> NOTE: For more information on generating a PagerDuty REST API key, see
> https://support.pagerduty.com/docs/generating-api-keys#generating-a-general-access-rest-api-key.

- **sl1_host**. The hostname or IP address of the SL1 system the incident states will synchronize to.
- **sl1_user**. The username for your SL1 system.
- **sl1_password**. The password for your SL1 system.
- **proxy_host**. The hostname or IP address of the HTTP proxy that will be used when making an API request to PagerDuty.
- **proxy_user**. The username for your associated HTTP proxy.
- **proxy_password**.The password for your associated HTTP proxy.
- **platform_identifier**. The parameter that is appended to the event ID to create the PagerDuty dedup key. When you configure your SL1 system to send event data to PagerDuty via PowerFlow, the "platform_identifier" parameter in the Run Book Action must match this value.
- **summary_string**. The string that reflects how the summery field is constructed before it is sent to PagerDuty. You can use Run Book Variables to dynamically include SL1 event data in the summary field. For example, the default value, "%S: %X: %M", includes the severity, affected device, and event message in the summary string.

> NOTE: For more information on Run Book Variables, see the "Run Book Variables" chapter in the *Run Book Automation* manual.

- **custom_details**. A structure of "key:value" pairs that can be used to provide additional SL1 event data to PagerDuty. You can use Run Book Variables to dynamically include SL1 event data in this field.
- **resolve_events**. A value to indicate whether events should be resolved in SL1 when the incident is resolved in PagerDuty. The default value is "False". If the value is "False", SL1 events will not be automatically cleared in SL1 when the corresponding PagerDuty incident is resolved.
- **link_events**. A value to indicate whether events should be updated with the incident URL from PagerDuty. The default value is "True". If the value is "True", the External Ticket ID field will be updated with the ID and URL of the PagerDuty incident.

> NOTE: SL1 supports only one value in this field. To avoid conflicts with other integrations, set this value to 'False' if your SL1 system has been integrated with another incident management system, e.g. ServiceNow.

6. Click **[Save]**. You can now align this configuration object with one or more applications.

Aligning a Configuration Object and Configuring PowerFlow Applications

In PowerFlow, you will need to align the PowerFlow applications for PagerDuty with the relevant configuration object, and, if needed, update any other fields on the **Configuration** pane for the applications.

To run this SyncPack, you must "align" the configuration object you just created with the following PowerFlow applications:

- ○ "Clear resolve event from ScienceLogic"
- ○ "Get PagerDuty log and incident data"
- ○ "Get resolve alert data from PagerDuty"
- ○ "Get trigger alert data from PagerDuty"
- ○ "Get alert and user data from PagerDuty and ScienceLogic"
- ○ "Post event updates to ScienceLogic"
- ○ "Create event in PagerDuty from SL1 event"

To align the configuration object with the relevant PowerFlow applications:

1. On the **Applications** page of the PowerFlow user interface, open one of the PowerFlow applications listed above and click **[Configure]**. The **Configurations** pane for that application appears.
2. From the **Configurations** drop-down, select the configuration object you want to use.

> NOTE: The values for **proxy_host** and the other parameters that appear in the **Configuration** pane with a padlock icon (🔒) are populated either by the configuration object you aligned with the application or by the Run Book Action. Do not modify these values. If you encounter an error, make sure your Run Book Action is configured properly.

3. Update any of the remaining fields on the **Configurations** pane as needed.
4. Click **[Save]**.
5. Repeat this process for the other PowerFlow applications.

Scheduling PowerFlow Applications

Using the "PagerDuty Incident Integration" PowerPack, you can optionally configure a single PowerFlow instance and single PagerDuty instance with multiple SL1 instances using the automation actions included in the PowerPack.

To run the automation, you must create a configuration object for each SL1 instance and schedule the "Get PagerDuty log and incident data" application.

You can create one or more schedules for a single application in the PowerFlow user interface. When creating each schedule, you can specify the queue and the configuration file for that application.

To create a schedule:

1. On the **Applications** page (▦), click the **[Schedule]** button for the application you want to schedule. The **Scheduler** window appears.

2. In the **Schedule List** pane, click the down arrow icon (⌄) next to an existing schedule to view the details for that schedule.

3. In the **Schedule Creator** pane, complete the following fields for the default *Frequency* setting:

   - *Schedule Name*. Type a name for the schedule.

   - *Frequency in seconds*. Type the number of seconds per interval that you want to run the application.

   - *Custom Parameters*. Type any JSON parameters you want to use for this schedule, such as information about a configuration file or mappings.

4. To use a cron expression, click the *Switch to Cron Expression* toggle to turn it blue. If you select this option, you can create complicated schedules based on minutes, hours, the day of the month, the month, and the day of the week:

As you update the cron expression, the **Schedule** window displays the results of the expression in more readable language, such as *Runs app: "Every 0 and 30th minute past every hour on Sat"*, based on 0,30 in the *Minutes* field and 6 in the *Day of Week* field.

5. Click **[Save Schedule]**. The schedule is added to the **Schedule List** pane. Also, on the **Applications** page, the **[Schedule]** button now displays with a dark blue background:
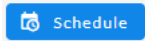
NOTE: After you create a schedule, it continues to run until you delete it. Also, you cannot edit an existing schedule, but you can delete it and create a similar schedule if needed.

To view or delete an existing schedule:

1. On the **Applications** page, click the **[Schedule]** button for the application that contains a schedule you want to delete. The **Scheduler** window appears.

2. Click the down arrow icon (⌄) to view the details of an existing schedule.

3. To delete the selected schedule, click the Actions icon (⋮) and select**[Delete]**.

TIP: On the **Scheduler** window for a PowerFlow application, you can click the **[Copy as]** button from the **Schedule List** pane to make a copy of an existing schedule.

NOTE: When either multiple SL1 instance or multiple PagerDuty instances are involved with PowerFlow, you should create an individual configuration object for each SL1 or PagerDuty instance. Next, create an individual schedule for each configuration object. Each schedule should use a configuration object that is specific to that single SL1 or PagerDuty instance. Creating copies of a PowerFlow application from a SyncPack for the purpose of distinguishing between domains is not supported, and will result in issues on upgrades.

ScienceLogic