



SL1 PowerFlow Platform Release Notes

Version 2.2.1

Overview

The *SL1 PowerFlow Platform* version 2.2.1 includes updates to the powerflowcontrol (pfctl) command-line utility and enhancements to the PowerFlow user interface. This release also addresses a number of issues.

NOTE: Unless mentioned elsewhere in the documentation, Synchronization PowerPacks do not require a specific version of the PowerFlow Platform.

NOTE: After the 2.1.0 platform release, the *Integration Service* was rebranded as *SL1 PowerFlow*, and the *Automation Builder* was rebranded as the *SL1 PowerFlow builder*.

This document covers the following topics:

<i>Features</i>	3
<i>Issues Addressed</i>	3
<i>Known Issues</i>	4
<i>System Requirements</i>	4
<i>Installing or Upgrading PowerFlow</i>	6

Features

The following features were included in *PowerFlow Platform* version 2.2.1:

The `powerflowcontrol` (pfctl) Command-line Utility

- The **healthcheck** action in the **powerflowcontrol** (pfctl) command-line utility was updated to display a detailed error message if the Couchbase version is not the same for all the nodes in a clustered environment.
- The **generate_haproxy_config** cluster-action in the **powerflowcontrol** (pfctl) utility creates an HAProxy configuration template that lets you easily set an HAProxy load balancer for a three-node cluster. For example: `pfctl cluster-action --action generate_haproxy_config`
- When configuring a three-node clustered environment, you can set the **OPEN_SECONDARY_CB_PORTS** configuration variable to "true" to expose Couchbase secondary ports through the main node IP or host name. You can set this configuration variable as a GUI environment variable in the **docker-compose.yml** file, or you can set it in the **isconfig.yml** file in the host. If **OPEN_SECONDARY_CB_PORTS** is set to "true", the GUI service exposes the Couchbase secondary ports in the compose file. The **autocluster** cluster-action in the **powerflowcontrol** (pfctl) utility was updated to automatically expose Couchbase secondary ports when creating a three-node clustered environment.

PowerFlow User Interface

- Added "Read Only" labels to tabs and views that cannot be edited in the PowerFlow user interface.
- Improved the error messaging when the **Trigger Application** operator does not specify a PowerFlow application to trigger.
- In the PowerFlow builder, when you drag a "flow control" operator from the **[Advanced]** tab to the canvas, the builder no longer automatically opens the corresponding wizard for that operator. This updated behavior lets you finish configuring the application before opening the wizard.
- On the **Transform Wizard** page in the PowerFlow builder, the variables, objects, functions, and scopes on the **[Jinja Template]** tab now display with different colors for easier reading and parsing.
- If PowerFlow retries a step in an application, PowerFlow no longer overwrites all logs of the previous run of that step. Instead, the logs from the retry are appended to the existing logs.

TIP: To view release notes and manuals for all versions of the SL1 PowerFlow Platform, see [SL1 PowerFlow Platform Documentation](#). To view release notes and manuals for PowerFlow Synchronization PowerPacks, see [SL1 Workflow Automation PowerPack Documentation](#).

Issues Addressed

The following issues were addressed in *PowerFlow Platform* version 2.2.1:

Synchronization PowerPack Installation

- Addressed an issue where users could not install and activate the Base Steps Synchronization PowerPack after upgrading from PowerFlow version 1.8.4 to version 2.1.x. (Case: 00130981 . JIRA IDs: EM-40288)
- For the "Activate & Install Syncpacks" application, the **retry_max** parameter for the "Activate Syncpack" and "Install Syncpack" steps was set to 3 by default. The time between those retries is calculated randomly based on the number of retries. These settings prevent the steps from colliding with each other when the steps are run in a environment with a large number of syncpack_steprunners that are trying to install a SyncPack at the same time in their respective volumes. (Case: 00125429. JIRA ID: EM-39930)

PowerFlow Applications and Steps

- Added a new "Defining Retry Options for a Step" topic in the *Managing PowerFlow Applications* chapter of the **SL1 PowerFlow Platform** manual.

Known Issues

This release contains the following known issues:

- When attempting to upgrade PowerFlow to version 2.2.x, 2.3.x, or 2.4.x, the RabbitMQ user interface might become inaccessible due to an incorrect RabbitMQ version in the **docker-compose.yml** file. This issue is addressed in PowerFlow version 2.5.0, so ScienceLogic recommends that you upgrade to version 2.5.0.
- If your PowerFlow system uses self-signed certificates, you must manually accept the certificate before you can upload Synchronization PowerPacks. Go to **https://<IP address of PowerFlow>:3141/isadmin**, accept the certificate, and then log into PowerFlow. After you log in, you will be able to upload Synchronization PowerPacks.
- The *latest* tag does not exist after the initial ISO installation. This situation only affects users with custom services that point to the *latest* tag. To work around this issue, run the tag latest script manually after running the `./pull_start_iservices.sh` command:

```
python /opt/iservices/scripts/system_updates/tag_latest.py  
/opt/iservices/scripts/docker-compose.yml
```

System Requirements

NOTE: PowerFlow Platform version 2.2.1 and later requires version 1.3.1 or later of the *Base Steps Synchronization PowerPack*. This version includes an update to the "Query REST" step that prevents issues with scheduled PowerFlow applications. You can download the latest version of this Synchronization PowerPack from the [PowerPacks](#) page of the ScienceLogic Support Site.

NOTE: The PowerFlow builder is available as part of an SL1 Premium solution. To upgrade, contact ScienceLogic Customer Support. For more information, see <https://sciencelogic.com/pricing>.

The PowerFlow platform does not have a specific minimum required version for SL1. However, certain PowerFlow Synchronization PowerPacks have minimum version dependencies. Please see the documentation for those Synchronization PowerPacks for more information on those dependencies.

The following table lists the port access required by PowerFlow:

Source IP	PowerFlow Destination	PowerFlow Source Port	Destination Port	Requirement
PowerFlow	SL1 API	Any	TCP 443	SL1 API Access
SL1 Run Book Action	PowerFlow	Any	TCP 443	Send SL1 data to PowerFlow
Devpi	PowerFlow	Any	TCP 3141	Internal Python package repository for Synchronization PowerPacks; check for self-certification for PowerFlow
Dex Server	PowerFlow	Any	TCP 5556	Enable authentication for PowerFlow
PowerFlow	SL1 Database	Any	TCP 7706	SL1 Database Access
powerflowcontrol (pfctl, formerly called iservicecontrol) command-line utility	PowerFlow	Any	22 (on all host nodes)	Log in and perform admin tasks on nodes
Encapsulated Security Protocol (ESP)	PowerFlow	IP Protocol 50	n/a	Security; ESP should be open and available between cluster nodes
Couchbase Dashboard	PowerFlow	8091	n/a	Couchbase Dashboard
RabbitMQ Dashboard	PowerFlow	15672	n/a	RabbitMQ Dashboard

ScienceLogic highly recommends that you disable all firewall session-limiting policies. Firewalls will drop HTTPS requests, which results in data loss.

CAUTION: PowerFlow clusters do not support vMotion or snapshots while the cluster is running. Performing a vMotion or snapshot on a running PowerFlow cluster will cause network interrupts between nodes, and will render clusters inoperable.

CAUTION: The site administrator is responsible for configuring the host, hardware, and virtualization configuration for the PowerFlow server or cluster. If you are running a cluster in a VMware environment, be sure to install open-vm-tools and disable vMotion.

CAUTION: You can configure one or more SL1 systems to use PowerFlow to sync with a *single* instance of a third-party application like ServiceNow, Restorepoint, or Cherwell. You cannot configure one SL1 system to use PowerFlow to sync with *multiple* instances of a third-party application like ServiceNow or Cherwell. The relationship between SL1 and the third-party application can be either one-to-one or many-to-one, but not one-to-many.

NOTE: You can use a single PowerFlow system to manage multiple pairings between one or more SL1 systems and third-party applications like ServiceNow and Cherwell. The pairings must always be one-to-one or many-to-one: one or more SL1 systems connected to only one third-party application.

NOTE: The default internal network used by PowerFlow services is **172.21.0.1/16**. Please ensure that this range does not conflict with any other IP addresses on your network. If needed, you can change this subnet in the **docker-compose.yml** file.

NOTE: The PowerFlow operating system is an Oracle Linux distribution, and all patches are provided within the standard Oracle Linux repositories. The patches are not provided by ScienceLogic.

TIP: For more information about system requirements for your PowerFlow environment, see the [System Requirements](#) page at the ScienceLogic Support site.

Installing or Upgrading PowerFlow

For detailed steps about installing or upgrading to this version of PowerFlow, see the *Installing and Configuring PowerFlow* chapter in the **SL1 PowerFlow Platform** manual.

© 2003 - 2022, 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/company/legal>.

ScienceLogic

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

International: +1-703-354-1010