



---

# SL1 PowerFlow Platform Release Notes

Version 2.6.0

---

# Overview

SL1 PowerFlow Platform version 2.6.0 includes updates to Couchbase, RabbitMQ, and Celery; updates to the **powerflowcontrol** (pfctl) utility; enhancements to the PowerFlow user interface; a new OpenTelemetry PowerFlow application; and other improvements.

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

This document covers the following topics:

<i>Features</i>	3
<i>Issues Addressed</i>	6
<i>Known Issues</i>	6
<i>System Requirements</i>	7
<i>Installing or Upgrading PowerFlow</i>	9

---

## Features

This section covers the features that are included in SL1 PowerFlow Platform version 2.6.0.

### Updates to Couchbase, RabbitMQ, and Celery

- This release of the PowerFlow Platform includes Couchbase version 6.6.0, which is a substantial upgrade from Couchbase version 6.0.2. As a result, please review the considerations in the "Upgrading to Couchbase Version 6.6.0" topic in the **PowerFlow Platform** manual.
- To troubleshoot issues with Couchbase or RabbitMQ, you can temporarily expose Couchbase or RabbitMQ ports without requiring Dex authentication. For more information, see the "Troubleshooting" chapter in the **PowerFlow Platform** manual.
- Added a new configuration for Celery to the **docker-compose** file that lets you send a task ID to the broker and load the task data from redis. You can set this new configuration by adding the **broker\_load\_from\_backend** environment variable. If you set this variable to True, Celery will save the payload it would have sent through the broker (RabbitMQ, SQS) and sends it to the backend (redis) instead, reducing the size of the message. When set to False, Celery will work as it normally does. The default value for this setting is False.

### Updates to the PowerFlow User Interface

- Added a set of filters for the columns for the inventory pages (pages that display lists of items) in the PowerFlow user interface. You can start typing filter text or select filter options in one or more of these filters to narrow down the items in the list to view just the items you want to view.
- On the detail page for a SyncPack on the **SyncPacks** page, the **Package Contents** section was updated to display more details about the steps, applications, and configuration objects include with that SyncPack:
  - The **[Applications]** tab includes a link to the detail page for each application, along with a list of steps in the application (with the option of clicking the step name to view the code for that step), and a thumbnail image of the workflow for the application.
  - The **[Steps]** tab includes a link to the Step Code dialog for each step, along with a list of any applications that are using that step. You can also click the name of the application from the list to go to the detail page for that application.
  - The **[Configurations]** tab includes a link to the detail page for each configuration object, along with a list of any PowerFlow applications that are aligned with that configuration object (with the option of clicking the application name to view the detail page for that application).
- Added an "Auto Refresh" counter that lets you see when data will update to the following pages: **SyncPacks**, **Applications**, **Configurations**, and **API Keys**. On these pages, you can also click "Auto Refresh" to refresh the page immediately.
- Pop-up messages have a countdown timer that displays until it closes, along with a Close icon.
- When you are scheduling a application, the name of that application now appears at the top left of the Scheduler window.

## Updates to the PowerFlow Control Tower Page

- Added RabbitMQ data and charts to the **System Health** widget.
- Added the **Connection\_widget** field to the "PowerFlow Check Connections" application in the "System Utils" SyncPack, which is included with this release of PowerFlow. You can use the new "steps" key in the **Connection\_widget** field to filter the data that displays on the **Workflow Health and Interconnectivity** widget on the **PowerFlow Control Tower** page:
  - When this field is present, the application filters for runs that used the "test" configuration object.
  - When this field is blank, the application filters for apps that did not have a configuration object aligned.
  - When the field is left out, the application does not filter, and it fetches and processes all runs regardless of the configuration objects (which is how the application worked by default in previous versions).

For example, in the **Connection\_widget** field, you can add the following JSON code to display the latest application run with the "test" configuration object, and the application will only show successful or failed runs with that configuration object:

```
{
  "name": "Integration Template",
  "steps": [
    {
      "app_name": "integration_template",
      "step_name": "Get REST Test",
      "configuration": "test",
      "syncpack": "base_steps_syncpack"
    }
  ]
}
```

## Updates to Configuration Objects

- In the "IS - System Backup Configuration Example" configuration object, the default value of the **backup\_destination** field was set to `/usr/tmp`.
- The **Configurations** page now lists configuration objects with their "user-friendly" names instead of the system IDs used by the API.

## Updates to the powerflowcontrol (pfctl) Command-line Utility

The latest version of the **powerflowcontrol** (pfctl) command-line utility, version 2.7.6 includes the following updates:

- Updated the pfctl utility so that only one instance of the utility can be run at the same time in the same stack.
- Improved error handling when the powerflowcontrol (pfctl) utility encounters a timeout when querying the pypiserver.
- The **healthcheck** action can detect duplicate versions of the "Activate Syncpack" and "Install Syncpack" steps, and the **autoheal** action deletes the duplicates.
- The **healthcheck** action checks the **/var** and **/tmp** directories on all nodes and raises an alert if either of the directories are at more than 91% capacity. (Case: 00303742)
- The **healthcheck** action can detect if any RabbitMQ queues are out of sync or missing a policy, and the **autoheal** action will sync those queues. These actions also detect and fix custom queues. If the actions cannot fix the queues automatically, the pfctl utility will display the steps to address the issue manually.
- Added a cluster action to the **healthcheck** action that gathers data from docker stats on each node.
- The **autocluster** action was updated to fix additional issues, including issues with the **docker-compose** file.
- You can check the installed pfctl version by running: `pfctl --version`. You can also view the current PowerFlow version and the installed pfctl version if you add `--json` at the start of the **healthcheck** command.
- Added additional **autoheal** and **create\_node\_indexes** steps to the "Troubleshooting Clustering and Node Failover" topic in the *Troubleshooting SL1 PowerFlow* chapter of the **SL1 PowerFlow Platform** manual.

**NOTE:** Do not use version 2.7.5 of the **powerflowcontrol** (pfctl) command-line utility. To download and install version 2.7.6, see [Installing the powerflowcontrol \(pfctl\) utility](#).

## Additional Features

- Added the "Collect PowerFlow OTEL Metrics" application to the "PowerFlow Check Connections" application in the "System Utils" SyncPack, which is included with this release of PowerFlow. This application collects basic PowerFlow OpenTelemetry content metrics, including:
  - The number of applications, steps, configuration objects, and SyncPacks.
  - The number of applications run during the last day.
  - Data about all of the content in the current PowerFlow system, including names and versions of all steps, configuration objects, and SyncPacks.
  - Platform services versions.

You can use this application to send the collected metrics and PowerFlow content to a metrics server in JSON format. You can also use this application to send those metrics to a OTEL collector using a OTEL GRPC endpoint, which is configurable as an application variable. For more information, see the **SL1 PowerFlow** manual. These features are disabled by default.
- You can run a **GET api/v1/status** operation that retrieves all the health status cache documents without running the "PowerFlow Control Tower HealthCheck" application.

- Updated the **GET /api/v1/status** operation to give you the option of retrieving all health metrics for a specific service, including the following services: contentapi, couchbase, dexserver, iservices\_syncpack\_steprunner, iservices\_syncpacks\_steprunner, pfctl\_output, rabbitmq, redis, steprunner.

For example: **GET https://<PowerFlow\_site>/api/v1/status/redis**.

- The following services are included in this release of PowerFlow:
  - **contentapi**. image: registry.scilo.tools/sciencelogic/pf-api:rhel2.6.0
  - **couchbase**. image: image: registry.scilo.tools/sciencelogic/pf-couchbase:6.6.0-9
  - **dexserver**. image: registry.scilo.tools/sciencelogic/pf-dex:2.22.0-5
  - **flower**. image: registry.scilo.tools/sciencelogic/pf-worker:rhel2.6.0
  - **gui**. image: registry.scilo.tools/sciencelogic/pf-gui:2.6.0-rc1-ubi7
  - **pypiserver**. image: registry.scilo.tools/sciencelogic/pf-pypi:6.3.1-9
  - **rabbitmq**. image: registry.scilo.tools/sciencelogic/pf-rabbit:3.8.35-3
  - **redis**. image: registry.scilo.tools/sciencelogic/pf-redis:6.2.7-2
  - **scheduler**. image: registry.scilo.tools/sciencelogic/pf-worker:rhel2.6.0
  - **steprunner**. image: registry.scilo.tools/sciencelogic/pf-worker:rhel2.6.0
  - **syncpacks\_steprunner**. image: registry.scilo.tools/sciencelogic/pf-worker:rhel2.6.0

---

## Issues Addressed

The following issues were addressed in this release:

- Addressed an issue where the powerflowcontrol (pfctl) version was downgraded during a PowerFlow upgrade. (Case: 00328661) (Jira ID: INT-5473)
- Addressed an issue where the **healthcheck** action generated false positives when **LOAD\_BALANCER: yes** is set in the **isconfig.yml** file for a three-node cluster. (Case: 00344031) (Jira ID: INT-4169)
- A number of updates were made to address and prevent potential security vulnerabilities. (Jira IDs: EM-58385, EM-58387, EM-58388, EM-58389)

---

## Known Issues

This release contains the following known issues:

- When upgrading to Couchbase version 6.6.0, the number of documents in the logs bucket could make the upgrade take longer, as a namespace upgrade is needed. ScienceLogic recommends that you flush the logs bucket if there are more than 300,000 documents that are taking up close to 2 GB of space in every node. Flushing the logs bucket will speed up the upgrade process. Otherwise, migrating a logs bucket of that size would take two to three minutes per node.

Run the following command to flush the logs bucket after the PowerFlow version 2.6.0 RPM was installed, but before redeploying the PowerFlow Stack:

```
pfctl --host <hostname> <username>:<password> node-action --action
flush_logs_bucket
```

Alternately, you can flush the logs bucket manually using the Couchbase user interface.

- For upgrades from PowerFlow version 2.2.x systems that have the **localpkg\_gpgcheck=1** option enabled in **/etc/yum.conf**, the SL RPM Public Key is required. Please contact your ScienceLogic Customer Success Manager (CSM) or create a new Service Request case at <https://support.sciencelogic.com/s> in the "PowerFlow" category to request access to that key.
- To avoid authentication issues, do not use the dollar sign (\$) character as the first character in any of the passwords related to PowerFlow. You can use the \$ character elsewhere in the password if needed.
- In PowerFlow version 2.4.0 and later, if you enabled the latest authentication updates for the backend services, the RabbitMQ API is no longer available externally from the cluster. As a result, remote API requests directly to RabbitMQ might not work (the RabbitMQ user interface is still completely operational). As a workaround, if you require remote access to the RabbitMQ API, you can return to legacy behavior by setting the following **gui** environment variable: **force\_auth\_validation: true**. Alternatively, you may perform any api requests to rabbit directly from within the container. Remote RabbitMQ API access for internal authentication users will be enabled in a future release of PowerFlow.
- The **Workflow Health and Interconnectivity** widget on the **PowerFlow Control Tower** page displays diagrams for PowerFlow applications and SyncPacks that have been deleted. To work around this issue, run the "PowerFlow Control Tower HealthCheck" application or wait for the next scheduled run of the application.
- If your PowerFlow system uses self-signed certificates, you will need to manually accept the certificate before you can upload SyncPacks. 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 SyncPacks.
- 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

PowerFlow Platform version 2.2.1 and later requires version 1.3.1 or later of the "Base Steps" SyncPack. 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 SyncPack 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 or AP2. However, certain SyncPacks for PowerFlow have minimum version dependencies, which are listed on the [Dependencies for SL1 PowerFlow SyncPacks](#) page.

## Ports

The following table lists the PowerFlow ingress requirements:

Source	Port	Purpose
SL1 host	443	SL1 run book actions and connections to PowerFlow
User client	3141	Devpi access
User client	443	PowerFlow API
User client	5556	Dex Server: enable authentication for PowerFlow
User client	8091	Couchbase Dashboard
User client	15672	RabbitMQ Dashboard
User client	22	SSH access

The following table lists the PowerFlow egress requirements:

Destination	Port	Purpose
SL1 host	7706	Connecting PowerFlow to SL1 Database Server
SL1 host	443	Connecting PowerFlow to SL1 API

## Additional Considerations

Review the following list of considerations and settings before installing PowerFlow:

- ScienceLogic highly recommends that you disable all firewall session-limiting policies. Firewalls will drop HTTPS requests, which results in data loss.
- Starting with PowerFlow version 3.0.0, the minimum storage size for the initial partitions is 60 GB. Anything less will cause the automated installation to stop and wait for user input. You can use the tmux application to navigate to the other panes and view the logs. In addition, at 100 GB and above, **PowerFlow** will no longer allocate all of the storage space, so you will need to allocate the rest of the space based on your specific needs.
- 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.
- 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.



- You can configure one or more SL1 systems to use PowerFlow to sync with a *single* instance of a third-party application like ServiceNow 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.
- 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.

For more information about system requirements for your PowerFlow environment, see the System Requirements page at the ScienceLogic Support site at <https://support.sciencelogic.com/s/system-requirements>.

---

## Installing or Upgrading PowerFlow

For detailed steps about installing or upgrading to this version of PowerFlow, see [Installing and Configuring PowerFlow](#).

**IMPORTANT:** You should always upgrade to the most recent release of PowerFlow.

© 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](mailto:legal@sciencelogic.com). For more information, see <https://sciencelogic.com/company/legal>.



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

International: +1-703-354-1010