Integration Service Platform Release Notes

Version 2.0.0, rev1

Overview

The Integration Service platform version 2.0.0 includes an updated user interface that matches the new user interface for SL1. The new user interface includes a new **Dashboards** page to help you monitor the resources used by the Integration Service, a new **SyncPacks** page that lets you import, install, and edit SyncPacks, and a new **Admin Panel** page for defining access permissions and user groups. You can also create integration applications and steps on the **Integrations** page. In addition, this release includes new authentication options, additional new features, and addressed issues.

WARNING: If you have made any customizations to default integration applications or steps that shipped with previous versions of the Integration Service, you will need to make those customizations compatible with Python 3.6 or later *before* upgrading to version 2.0.0 of the Integration Service. Also, if you made any modifications to the nginx configuration or to other service configuration files outside of the **docker-compose.yml** file, you will need to modify those custom configurations before upgrading, or contact ScienceLogic Support to prevent the loss of those modifications. Please note that upgrading to version 2.0.0 *will* involve some downtime of the Integration Service.

This document describes:

- *Pre-install or pre-upgrade information*
- *The upgrade process*
- *The features included in version 2.0.0*
- *The issues addressed in version 2.0.0*
- *The known issues in version 2.0.0*

Before You Install or Upgrade

The Integration Service platform has the following minimum system requirements. Please note that these system requirements ultimately depend on the amount of workload you plan on running on your Integration Service:

- 8 CPUs
- 24 GB total RAM
- 100 GB total storage

NOTE: The Integration Service needs its own dedicated memory. Thin provisioning is not supported.

The following table offers a conservative starting point for sizing based on a typical environment (any object being processed by the Integration Service is considered a synced object):

Minimum at 1,000 Synced Objects			Minimum at 10,000 Synced Objects			Minimum at 50,000 Synced Objects		
RAM (GB)	Cores	Disk (GB)	RAM (GB)	Cores	Disk (GB)	RAM (GB)	Cores	Disk (GB)
24	8	100	36	8	100	48	8	200

All workloads are different. Storage requirements will vary based upon monitoring depth, frequency of integrations, and length of retention. Sizing recommendations may differ based on multi-SL1 stack support.

NOTE: The Integration Service itself does not have a specific minimum required versions for SL1 or AP2. However, certain Integration ServiceSynchronization PowerPack have minimum version dependencies. Please see the documentation for those Synchronization PowerPack for more information on those dependencies.

The following table lists the port access required by the Integration Service:

Source IP	Integration Service Destination	Integration Service Source Port	Destination Port	Requirement
Integration Service	SL1 API	Any	TCP 443	SL1 API Access
SL1 Run Book Action	Integration Service	Any	TCP 443	Send SL1 data to the Integration Service
Devpi	Integration Service	Any	TCP 3141	Internal Python package repository for Synchronization PowerPacks; check for self-certification for the Integration Service
Dex Server	Integration Service	Any	TCP 5556	Enable authentication for the Integration Service
Integration Service	SL1 Database	Any	TCP 7706	SL1 Database Access
Docker Visualizer	Integration Service	8081	n/a	Docker Visualizer (**http://** only)
Couchbase Dashboard	Integration Service	8091	n/a	Couchbase Dashboard (use your Integration Service credentials)
RabbitMQ Dashboard	Integration Service	15672	n/a	Rabbit MQ Dashboard (use *guest/guest*for credentials)

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

NOTE: The Integration Service 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.

Upgrade Process

This topic explains at a high level how to upgrade to Integration Service version 2.0.0. The process for upgrading to version 2.0.0 of the Integration Service includes the following required steps:

1. Upgrade the host packages and Python 3.6 (previous versions of the Integration Service used Python 2.6).
2. Upgrade to Oracle 7.3 or later.
3. Upgrade to Docker version 18.09.2 or later.

> **NOTE:** Integration Service version 2.0.0 or later requires the **docker-ce 18.09.2** or later version of Docker. The Integration Service ISO installs the **docker-ce 19.03.5** version of Docker by default, but if you are upgrading to this version from the RPM, you must first upgrade Docker before you upgrade the Integration Service with the RPM.

4. Install the Integration Service 2.0.0 upgrade RPM.
5. Update the Integration Service system from Basic Authentication to OAuth 2.0.
6. Set up licensing for the Integration Service. After installation, you must license your Integration Service system to enable all of the features. For more information, see the "Managing Licenses for the Integration Service" chapter in the ***Integration Service Platform*** manual.

Please note that upgrading to version 2.0.0 *will* involve some downtime of the Integration Service.

You can perform the upgrade steps manually, or you can run the **is_upgrade_to_2.0.0.sh** script to perform the upgrade steps automatically. The script upgrades the Integration Service from 1.8.x to 2.0.0.

To locate the upgrade script:

1. Go to the ScienceLogic Support site at https://support.sciencelogic.com/s/.
2. Click the **Product Downloads** tab and select *Integration Service*. The **Integration Service Release** page appears.
3. Click the "Integration Service 2.0" link. The **Integration Service 2.0 Release Version** page appears.
4. In the **Release Files** section, click the "1.8.X to 2.0.0 Upgrade" link for the script and click Download File. The **is_upgrade_to_2.0.0.sh** script is in the **is_upgrade_tools.zip** file.

For detailed steps about *upgrading* the Integration Service, see the "Upgrading the Integration Service" topic in the "Installing and Configuring the Integration Service" chapter in the ***Integration Service Platform*** manual.

For detailed steps about *installing* the Integration Service, see the "Installing and Configuring the Integration Service" chapter in the ***Integration Service Platform*** manual.

Included Features

Integration Service version 2.0.0 includes the following features:

General

- The Integration Service user interface was updated to match the SL1 user interface, with the navigation tabs now located on the left-hand side of the window. The tabs provide access to the following pages: **Dashboard**, **SyncPacks**, **Integrations**, **Configurations**, and **Reports**. Also, a number of text buttons were replaced with icon buttons, and the user interface was updated to display human-readable names and labels without underscores. For more information, see the ***Integration Service Platform*** manual.
- The **About** page (username > About) of the Integration Service user interface includes the version numbers for the API Package, the Common Package, and the GUI, along with information about the license used by the Integration Service.
- In the new **Notification Center** pane (username > Notifications) you can view a log of all previous notifications that appeared on the Integration Service system. From the **Notification Center** pane, you can click links that appear in notifications and go to that run of an integration application.
- Security for the Integration Service was enhanced. If a user does not enter a password after a certain amount of time, the SSH session with the Integration Service server times out. Also, SSH sessions will time out after a certain amount of time, unless overridden by user.
- You can apply the new **log_rollover** environment variable to steprunner services to control the number of log file backups that are present when logging to files. The default setting for this variable is five backups.
- While the **SyncPacks**, **Integrations**, **Configurations**, **Reports**, and **Admin Panel** pages are loading, a dark green, horizontal line appears at the top of the page until the loading is complete.

Dashboard page

- The new **Dashboard** page in the Integration Service user interface displays a list of the active, processed, failed, succeeded, and retried tasks on that Integration Service instance. This page is now the initial landing page after you log in, and it is the default page when navigating to the instance.
- The **Dashboard** page includes interactive graphs that display task status, worker status, and application status.

SyncPacks page

- The new **SyncPacks** page in the Integration Service user interface lets you import and install Synchronization PowerPacks. By default, the **SyncPacks** page displays all activated SyncPacks. You can click the Show All SyncPacks toggle at the top of the **SyncPacks** page to view all versions of all available SyncPacks.
- If you click the name of a SyncPack on the **SyncPacks** page, you can view more information about the SyncPack package and its contents, include the integration applications, configurations, and steps included in that SyncPack.
- You can access the logs related to SyncPack installation and activation by clicking the icon in the **Activation** column for that SyncPack on the **SyncPacks** page.

- SyncPack developers can create new SyncPacks that contain a custom set of integration applications, configuration objects, and steps. Developers can then export and share that SyncPack with other Integration Service users. For more information, see the ***Integration Service for Developers*** manual.
- SyncPack developers can define step parameters as constants and use those parameters in multiple applications and steps.
- All ServiceNow-specific content, including integration applications, steps, and configuration objects, were moved from the Integration Service platform and into separate ServiceNow SyncPacks.
- The existing base steps from previous versions of the Integration Service were added to the *ServiceNow Base Steps SyncPack*. This Synchronization PowerPack is added to the **SyncPacks** page when you install or upgrade to Integration Service version 2.0.0.
- The ***Integration Service Developers Guide*** version 2.0.0 was updated with new information about creating content for the Integration Service, including creating SyncPacks, integration applications, and steps; editing the components in a SyncPack, and adding custom libraries to a SyncPack.

Integrations page

- On the **Integrations** page, you can create new integration applications and new steps for applications. You can also edit and delete existing integration applications, and you can edit and delete the code for existing steps.
- The **SyncPack** column on the **Integrations** page lists the name of the SyncPack to which a specific integration application belongs, where relevant.
- You have the following options when running an integration application from an **Integration Application** page:
    - ***Run***. Executes the integration application normally, with a log level of 1. This is the default.
    - ***Debug Run***. Executes the integration application in Debug Mode with a log level of 10.
    - ***Custom Run***. Executes the integration application using a logging level that you specify, along with any customer parameters you want to use to test specific features in the application.
- The backup and restore integration applications were updated to enable the backup of SyncPack data. The "Integration Service Backup" integration application includes two new fields on the **Configuration** pane: ***include_syncpacks*** and ***installed_only***. The "Integration Service Backup" integration application includes two new fields: ***include_syncpacks*** and ***force_syncpack_upload***. Also, the "is-system-backup-configuration-example" configuration object was added to the System Utils SyncPack along with the update backup and restore integration applications.
- The ***Integration Service Developers Guide*** version 2.0.0 contains information about the ability to specify the queues where steps can execute when running an integration application.
- The **Steps Registry** pane on the **Integrations** page was updated to show all steps available on the Integration Service server. You can click the Actions button on a step in the registry to view more information about that step and to edit the step code.
- You can add mock data to a step in an integration application and click **Custom Run** to run test data for that step.
- To view a pop-up description of a field on the **Configuration** pane for an integration application, hover over the label name for that field.

- When you run an integration application and a warning occurs in the step logs for that application, the box around the step with the warning will be outlined in yellow to draw attention to it.

Configurations page

- The **Configurations** page includes an Actions button that lets you edit or delete a specific configuration object.
- When creating or editing a configuration object on the **Configurations** page, the name-value pairs in the Configuration Data section display in text fields by default instead of a block of JSON code. For more complex configuration objects, you can use the Toggle JSON Editor button to switch between text fields and JSON code for the configuration data.
- You can make copies of existing configuration objects to simplify deployment.

Reports page

- You can view reports from the "Integration Service System Diagnostics" integration application on the new **Reports** page.

Admin Panel page

- You can integrate Integration Service version 2.0.0 with LDAP, Active Directory, SAML 2.0, and other authentication providers to support multiple users on the platform.
- On the new **Admin Panel** page, if you have the *Administrator* role, you can create one or more user groups to define the different roles for your users. Depending on their roles and permissions, users have access to certain features, or they are blocked from certain features. Roles include *Administrator*, *Develop*, *Configure*, *Execute*, and *View*, and these roles are defined in the ***Integration Service Platform*** manual.

API

- The new **/syncpacks** API endpoint lets you POST, GET, and DELETE SyncPack data, including installing SyncPacks. For more information, see the "API Endpoints" chapter in the ***Integration Service Platform*** manual.
- The new **/license** API endpoint lets you POST, GET, and DELETE licensing data.
- The new **/roles** and **/roles/owner** API endpoints lets you POST, GET, and DELETE roles assigned to users.
- If you want to make inbound requests to the Integration Service from an external system where customizing the POST payload is not possible, you can pass in parameters as part of an API request.

Installation and Configuration

- The process for upgrading to version 2.0.0 of the Integration Service includes a number of required steps. You can perform these upgrade steps manually, or you can run the **is_upgrade_to_2.0.0.sh** script to perform the upgrade steps automatically. An additional script is available for clustered environments. For more information, see the "Upgrading the Integration Service" topic in the "Installing and Configuring the Integration Service" chapter in the ***Integration Service Platform*** manual.
- To access this release, you need to create and upload a license key to your Integration Service instance. For more information, see the "Licensing the Integration Service" topic in the "Installing and Configuring" chapter of the ***Integration Service Platform*** manual.
- All common packages and SyncPacks run on Python 3.6. Python 2.7 is no longer supported.

- To support the new features on the **SyncPacks** page, a new Python Package Index (PyPI) service was added to the Integration Service stack. When deploying the Integration Service in a cluster setup, and not using network-aware volumes, the PyPI server must be "pinned" to a specific node with constraints. Pinning the PyPI server to a single node ensures that its persistent volume containing the SyncPacks will always be available to the Integration Service.
- All Dex and lua-resty-oidc module configuration for Integration Service authorization is included in the **isconfig.yml** file located at **/etc/iservices/**.
- Updated the **is_gui** image to deploy only one replica per node to let Integration Service administrators scale the user interface for the Integration Service for multiple nodes.
- Integration Service administrators can use a command-line utility called ***iservicecontrol*** that performs multiple administrator-level actions on either the node or the cluster. You can use this script to automate the configuration of a three-node cluster. For more information, see the "Configuring the Integration Service for High Availability" chapter of the ***Integration Service Platform*** manual.
- The ***Integration Service Platform*** manual was updated to explain how multi-tenant deployments can use separate virtual hosts and users for each tenant. For more information, see the "Advanced RabbitMQ Administration and Maintenance" topic in the appendix manual.
- The following services are included in the Integration Service version 2.0.0:
    - **steprunner**. image: sciencelogic/is-worker:2.0.0
    - **rabbitmq**. image: sciencelogic/is-rabbit:3.7.14-3
    - **flower**. image: sciencelogic/is-worker:2.0.0
    - **scheduler**. image: sciencelogic/is-worker:2.0.0
    - **redis**. image: sciencelogic/is-redis:4.0.11-2
    - **couchbase**. image: sciencelogic/is-couchbase:2.0.0
    - **contentapi**. image: sciencelogic/is-api:2.0.0
    - **visual**. image: dockersamples/visualizer
    - **gui**. image: sciencelogic/is-gui:2.0.0
    - **pypiserver**. image: sciencelogic/is-pypi:4.8.1-2
    - **dexserver**. image: sciencelogic/is-dex:2.18.0-1

Issues Addressed

The following addressed issues are included in Integration Service version 2.0.0:

- Addressed an issue where Docker containers were not able to write logs and the Integration Service stopped working because the **/root** and **/var/log** directories were at 100%.
- The Debug logging for the Integration Service Scheduler now provides additional information, including the integration application name being triggered, the parameters used by that application, and the schedule entry that started the run.

- Addressed an issue where the **sl1_user** value was case-sensitive when it should have been lowercase only.
- Addressed an issue where selecting *none* from the ***Configuration*** dropdown on the **Configuration** pane did not clear the selected configuration object from an integration application. Also, if you did not select a configuration object when editing fields on the **Configuration** pane, the previously set configuration object will remain aligned (if there was a previously set configuration object).
- Addressed an issue where scheduled integration applications were not properly honoring the ***loglevel*** application parameter.
- Addressed an issue where default values for some variables on the **Configuration** pane for an integration application were displayed as the actual values.
- Used hovertext to address an issue where long variable names on the **Configuration** pane were being truncated.
- Addressed an issue in the user interface with the "Sync Devices from SL1 to ServiceNow" integration application, where any values in the ***customer_ci_relation_overrides*** field were not saved or applied.
- Addressed an issue where the "Integration Service Backup" and the "Integration Service Restore" integration applications continued working in data-transfer mode after you set the ***data-only*** parameter to false.
- Removed a problematic index which could cause the indexer to end unexpectedly or cause high CPU usage because of abnormally large app_vars.
- Updated idx_logs_log_type_application_log with a critical fix to prevent the indexer from ending unexpectedly when it encountered abnormally large app_vars.
- Addressed an issue where special characters were not allowed in the display name for a configuration object.

Known Issues

- If your Integration Service system uses self-signed certificates, you will need to manually accept the certificate before you can upload SyncPacks. Go to **https://<*IP address of Integration Service*>:3141/isadmin**, accept the certificate, and then log into the Integration Service. After you log in, you will be able to upload SyncPacks.
- You must edit the code in the JSON Configuration editor if you want to add data to the **customer_ci_relation_overrides** application variable in the "Sync Devices from SL1 to ServiceNow" integration application. Click the Show JSON Configs button on the **Configuration** pane to access the JSON code.
- 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
  ```

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

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

International: +1-703-354-1010