



SL1 11.3.1 Release Notes

SL1 version 11.3.1 (Document revision 7)

SL1 Platform Release Notes, version 11.3.1

IMPORTANT: ScienceLogic strongly recommends that you review the [installation and upgrade instructions](#), [important notes about upgrading SL1](#), and [known issues for this release](#) before installing or upgrading to SL1 version 11.3.1.

The SL1 11.3.1 release restores the ability to license SL1 appliances using the classic Web Configuration Utility. It also includes security and package updates and numerous issue fixes.

This document covers the following topics:

New Features and Enhancements in 11.3.1	3
Issues Addressed in 11.3.1	3
New Features and Enhancements in 11.3.0	4
Issues Addressed in 11.3.0	14
Installing and Upgrading SL1	19
Important Upgrade Notes	20
Known Issues for SL1 Version 11.3.1	30
SL1 Extended Architecture	31
Recently Deprecated Features	32

New Features and Enhancements in 11.3.1

This section describes the features and enhancements that are included in SL1 version 11.3.1.

Licensing

- Restored the ability to license SL1 appliances using the classic Web Configuration Utility, which had been removed in SL1 11.2.0. The classic Web Configuration Utility can be accessed at `<IP address>:7700`.

PhoneHome

- In SL1 11.3.1, Military Unique Deployment (MUD) users are able to use the pre-11.2.0 PhoneHome services. With this change, MUD users are able to configure PhoneHome communications using the Web Configuration Utility that was previously used for that function prior to SL1 11.2.0 at `<IP address>:7700`. MUD users who use the pre-11.2.0 PhoneHome services cannot access the Node Configuration Utility that was introduced in 11.2.0 at `<IP address>:7700/node-config`; they will receive an error message if they attempt to do so.

Security

- 11.3.1 includes multiple package updates to improve security and system performance.

Issues Addressed in 11.3.1

This section describes the issues that were addressed in SL1 version 11.3.1.

API

- Ensured that the ScienceLogic API successfully returns IP addresses for network interfaces when a specific interface is queried through the API. (Jira ID: EM-57529)(Case: 00335451)

Authentication

- SL1 can now successfully validate signed SAML messages and assertions to enhance Active Directory Federation Services (ADFS) and Azure AD compatibility. If you are using ADFS, ensure that the **Integrated Windows Auth** field is set to *Enabled* on the **SSO Auth Resource Editor** page (System > Settings > Authentication > Resources > Actions > Create SSO Resource). (Jira ID: SLS-768)

Documentation

- Updated the **System Administration** manual chapter on "Upgrading SL1" to remove instructions relating to backing up and restoring the NextUI files. (Jira ID: EM-57749) (Case: 00338447)

Dynamic Applications

- Improved the deviation crunch process performance by no longer managing active policies for deviation crunch or permitting invalid minimum/maximum settings for collection objects that have deviation alerting enabled. (Jira IDs: EM-54388, EM-56905) (Case: 00291732)

User Accounts

- Resolved an issue in which users not assigned to an admin policy were unable to access all organizations and devices due to the following error: "GraphQL Error: Forbidden." (Jira ID: SLUI-16455, SLUI-15792)

User Interface

- Updated the **Agents** page (Devices > Agents) to ensure that it updates with new agent inventory rows as you scroll down the page. (Jira ID: SLUI-16887) (Case: 00327060)
- Updated the **[Registered]** tab of the **Nodes** page to increase the number of appliances that display to 25 and to automatically display the next 25 appliances when you scroll to the bottom of the list. (Jira ID: EM-56852) (Case: 00319543) (Case: 00320829)

New Features and Enhancements in 11.3.0

This section describes the features and enhancements that are included in SL1 version 11.3.0.

Agent

- You can now delete an Agent device on the **Device Manager** page (Devices > Device Manager) as well as on the **Agents** page (Devices > Agents). (Jira ID: AP-2012)
- Replaced the undocumented Windows API calls used by the SL1 Agent with Windows Performance Counters. (Jira ID: AP-2185)
- The Exclude function in the SL1 Agent can now read a configuration file encoded in UTF-16. (Jira ID: AP-2254)
- The SL1 Agent was updated to reduce the number of event log messages that display. (Jira ID: AP-893)
- Added the Agent installer Helm chart to the deployment for SL1 11.3.0 and later systems. As this chart is not deployed to SL1 systems before 11.3.0, this also updates the deployment of the ingress charts so that systems before 11.3.0 use the currently available ingress chart version. (Jira ID: AP-2177)
- Updated the Agent Windows Service alerting service to be more robust and to report as expected, including in situations where a service that no longer exists is aligned to a monitor. (Jira ID: AP-1848)

IMPORTANT: In SL1 version 11.3.0 and earlier, the SL1 Agent installs and runs as root. If you use the SL1 Agent, be advised that starting with the upcoming SL1 version 12.1.0, the agent will install and run as a new dedicated SL1 user account that is a non-root/non-administrator account by default. You will have the option to continue running the agent as root if you choose. Your

existing agents will continue to work in SL1 version 12.1.0, but you will need to update your current processes, automations, and any internal documentation related to the agent for the new user account. Additional information about this change will be in the release notes and related documentation for SL1 version 12.1.0.

Anomaly Detection

- The "Select Metric to Disable/Enable Machine Learning" prompt when bulk enabling or disabling metrics was renamed to "Select Available Metrics" or "Disable Following Metrics," respectively.
- A "Successfully enabled" notification was added that displays when a metric is enabled.
- An "Error enabling: \${error.message}" notification was added when a metric cannot be enabled due to an error occurring.
- The **[Enable]** button that you can select to bulk enable metrics was moved to the top of the **Machine Learning** page.
- Updated Anomaly Detection status to display "Building" or "Failed" when no historical data is available for a device metric. You can then try re-enabling the metric in the future to check if data is available.
- Updated the "Invalid data for stored prediction" notification so that it will not display when no predictions are available.
- Enhanced model creation speed.

Authentication

- For LDAP/AD accounts, a new **User Display Name** field was added to the **Account Permissions** page (Registry > Accounts > User Accounts > wrench icon > Permissions).
- For LDAP/AD accounts, a new **User Principal Name** field was added to the **Account Properties** page (Registry > Accounts > User Accounts > wrench icon > Properties).
- Two new fields were added to the **Authentication Resource Editor** page for LDAP/AD Authentication Resources (System > Settings > Authentication > Resources > Actions > Create LDAP/AD Resource):

- The **User Display Name** drop down field.
- The **UPN** field is a new addition for the User Principal Name.

If you are upgrading from a previous release, the **UPN** field will be blank for existing authentication resources. If you want to use the **User Display Name** field with the UPN for these existing resources, you must manually enter "userPrincipalName" in the **UPN** field. Newly created authentication resources will have that default UPN value included automatically.

Bandwidth Billing Policies

- The **Organization** field in the **Bandwidth Policy Editor** (Registry > Service Provider Utilities > Bandwidth Billing) is now editable.

Business Services

- You can now edit the **Poll Frequency Har Provider** value in the **Business Services** page's **[Info]** drop-down menu.
- The Business Services **[Overview]** tab now fetches device data only when you request that specific device service through the Sunburst or Overview table components. This results in faster load times.
- You can now designate services as favorites by clicking the star icon for those services from the **Business Services** page or from the **Overview** tab. Favorite services display at the top of the inventory list, alphabetically by default. In addition, you can filter and sort by favorite services on the **Business Services** page.

Collector Groups

- New columns were added to the **Collector Groups** page (Manage > Collector Groups), some of which are hidden by default. You can customize your column preferences by clicking Grid Settings > Column Preferences.
- New fields were added to the **Add/Edit Collector Group** modal (Manage > Collector Groups > Add Collector Group, or Manage > Collector Groups > actions > Edit).
- You cannot delete a collector group if it has any devices, data, or Message Collectors aligned to it. You also cannot delete a collector group if you do not have common organization alignment with the collector group. This means that:
 - You can create or edit collector groups that are aligned to all organizations only if you have access to all organizations.
 - You can align collector groups to all organizations only if you have access to all organizations.
 - If you have access to only specific organizations, you will be limited to aligning collector groups to only those organizations. You will not be able to affect collector groups' alignment with organizations to which you do not have access.
- Added failover validations to various fields.

Collector Pipeline

- Updated the Bitnami Redis container for Helm charts used for Collector Pipeline services.

Concurrent PowerShell Collection

- Updated the *SL1: Concurrent PowerShell Monitoring* PowerPack with several new and updated internal event policies and alerts for PowerShell connectivity and execution issues.

Credentials

- PowerShell credentials are now encrypted over HTTP rather than HTTPS if the credential's **Encrypted** field is set to *No*.

Data Retention

- You can now define data pruning thresholds for SSL certificates, ports, services, processes, and file systems. You can do this globally using new fields that were added to the Collection Data Retention section of the **Data Retention Settings** page (System > Settings > Data Retention), or you can define these settings at the device level on the **Device Thresholds** page (the **[Thresholds]** tab of the Device Investigator, or Registry > Devices > Device Manager > wrench icon > Thresholds in the classic SL1 user interface).
- Updated the Hourly Maintenance process to prune expired SSL certificate, port, service, process, and file system data from the database.

Deployment

- 11.3.0 includes an upgrade to Kubernetes 1.22.
- Military Unique Deployment (MUD) users can now deploy SL1 on AWS GovCloud.

Devices

- You can now designate devices as favorites by clicking the star icon for those devices from the **Devices** page or from the **Device Investigator Settings** page. Favorite devices display at the top of the inventory list, alphabetically by default. In addition, you can filter and sort by favorite devices on the **Devices** page.
- Added new **[Summary]** and **[Performance Metrics]** tabs to the **Device Investigator**. These tabs replicate the content of the **[Device Summary]** and **[Device Performance]** tabs from the classic SL1 user interface.
- Updated the **Collections** page in device details to add support for single row selection, which opens an existing sidebar.
- Added the "Support: Device Group Performance" Dynamic Application to the *ScienceLogic Support Pack PowerPack* to collect device group telemetry data when the Dynamic Application is aligned to any device in a device group.
- Updated the alert messages for the Device Maintenance windows to display details about the SL1 schedule that created the maintenance window, including the name of the schedule, its schedule ID, the duration, and whether the schedule applies to the device itself or a device group.

Discovery

- Updated the guided discovery wizard to enable you to go directly to a screen in the wizard by using a correctly-formatted URL.

Dynamic Applications

- Snippet and Internal Collections Dynamic Applications were updated to be compatible with Python 3. You can now write snippets for these Dynamic Application types using Python 3. (If you currently have any existing Dynamic Applications that use custom Python snippet code, see an important note in the section on [Libraries and Execution Environments](#).)
- In the **Dynamic Application Editor** page (System > Manage > Dynamic Applications > wrench icon), the **Execution Environment** field now displays the Python version in parentheses next to the environment name.

- Added target (self.device_ip, self.device_hostname, self.device_pdu_packing) and root (self.root_ip, self.root_hostname, self.root_pdu_packing) device attributes to the snippet context in snippet Dynamic Applications to better support collector affinity.
- When a Dynamic Application is in debug mode, logs are available in /var/log/sl1/sl_process.log to record choose_python_version functionality.
- Implemented a function that validates the format of results collected by Dynamic Applications. This function logs warning messages and stops execution before postprocessing when it encounters collected result objects that are formatted incorrectly.

Events

- Implemented a new Event Insights feature that collects alert metrics from the system, from both internal and external origins, and normalizes the data at 15 minute, 1 hour, and 24 hour intervals. This collection provides more focused metrics and context for the volumes of alerts handled by SL1, so you can see how SL1 uses this data to reduce the number of alerts that become events or tickets. The **Event Insights Overview** page includes the following chart widgets:
 - Two chart widgets, "Events Created" and "Active Events", that display trend lines representing a count of new and active event records created within a specified time period
 - A Sankey chart widget that displays event metric data over a specified time period
 - An Alert Counter line chart widget that displays alert data over a specified time period
- The "Maint: Dynamic App object collection suspended" event policy was renamed "Maint: Dynamic Application Collection Object Suspended", and its internal alert message has been modified slightly. In addition, a new "Maint: Dynamic Application Collection Object Enabled by Dynamic Refresh" event policy and internal alert were added. This event policy is disabled by default but is the inverse to the "Maint: Dynamic Application Collection Object Suspended" event policy.
- Added "Component unmerged with device" and "Component unmerge with device failed" event policies to the *SL1 Default Internal Events PowerPack*. You can enable these policies to trigger when a component is successfully or unsuccessfully unmerged from a device.
- SL1 generates an event for incorrect Windows PowerShell passwords.
- PowerShell communication events that were generated in SL1 will be cleared as soon as the affected device becomes available.

GraphQL

- Introduced a new "topology" GraphQL resource for Business Services to replace "relatedNodes."
- Added GraphQL support for event and alert metric data in Sankey chart widgets on the Event Insights Overview page.
- You can now modify additional fields with the "createCollectorGroup" and "updateCollectorGroup" GraphQL mutations.
- In GraphQL, you can toggle the "allOrgs" field for a "collectorGroup" resource only if you have All Organizations access.

- Added new fields relating to failover, fallback, virtual collector group, and concurrent collections settings that can be queried in the "collectorGroup" GraphQL resource.
- Added two new GQL resources, "Schedule" and "Task", to support new scheduler features in SL1.
- Added the following GraphQL queries for Setup & Configuration workflows and activities: setupConfigWorkflow, setupConfigWorkflows, setupConfigActivity, and setupConfigActivities.
- Added a new "updateSetupConfigActivityStatuses" GraphQL mutation that enables you to bulk update the statuses for multiple Setup & Configuration activities.

Installation

- SL1 no longer specifies the default credentials when adding an SL1 collector to a stack. As a result, you must specify the credentials for the collector on the **Appliances** page (System > Settings > Appliances) in the **DB User** and **DB Password** fields.
- If an appliance is added with the wrong appliance type, SL1 generates a critical error to notify you.
- When a new appliance is created, upon registration or approval of the request, the PhoneHome server encrypts the passwords before storing them in the licenses table.

Libraries and Execution Environments

- You can now write Dynamic Application snippets using Python 3, create Python 3 execution environments, align Python 3 execution environments to Run Book Actions, and import Python 3 ScienceLogic Libraries.
- Newly created execution environments will use Python 3 by default.
- For existing execution environments, you can now use a new **Environment Type** field on the **Environment Editor** page (System > Customize > ScienceLogic Libraries > Actions > Execution Environments > wrench icon) to edit the Python version.
- A new **Env Type** column now appears on the **Environment Manager** page (System > Customize > ScienceLogic Libraries > Actions > Execution Environments) to indicate the Python version for each execution environment. You can use this column to sort the list by Python version or filter the list for specific Python versions.
- On the **ScienceLogic Library Manager** page (System > Customize > ScienceLogic Libraries), a new **Python Version** column indicates the Python version of each library listed. You can use this column to sort the list by Python version or filter the list for specific Python versions.
- When creating or editing an execution environment that has aligned ScienceLogic Libraries, you can view but not edit the environment type. This prevents a scenario where ScienceLogic Libraries are using a different Python version than the execution environments to which they are aligned.
- When importing a ScienceLogic Library, if the library has an improper "requires_python" value, the library will not be imported and an error message will be displayed.

IMPORTANT: Prior to SL1 11.3.0, all Dynamic Application snippets, Execution Environments, Run Book Actions, and ScienceLogic Libraries utilized Python 2. With the introduction of Python 3 support in this release, ScienceLogic intends to deprecate support for Python 2 in a future release. At that time, any custom Python code that you have written within SL1 will cease to

work properly. Therefore, if you currently use custom Python 2 code, ScienceLogic strongly recommends that you proactively convert it to use Python 3 instead. Additional information about this change will be in the release notes and related documentation for the SL1 version in which Python 2 support is deprecated.

Licensing

- Removed all licensing information and licensing menu options from the classic Web Configuration Utility at `<IP address>:7700`. You can revert back to the classic Web Configuration Utility by clicking the "Revert back to classic" banner displayed on the Node Configuration Utility.

NOTE: The classic Web Configuration Utility *was restored in SL1 11.3.1*.

- The default SL1 license capacity was updated to 1,000 devices for 90 days.

Logging

- ScienceLogic logging functions have been updated to log using rsyslog instead of writing directly to files.
- Configuration files for rsyslog were completely updated. This new configuration gives you the option of configuring TLS to send or receive syslog messages, forwarding logs to a security information and event management (SIEM) tool, filtering inbound logs, and other features.

WARNING: Any existing modifications you made to your rsyslog configurations to support log forwarding, filtering, or TLS reception before SL1 version 11.3.0 will be removed. To reconfigure any custom rules using the appropriate syntax, see the "Logging in SL1 Version 11.3.0 and Later" topic in the *Daily Health Tasks* section of the *System Administration* manual.

- Updated the Audit Log entries to include the appliance ID for one-time password requests.
- Added new logging functionality to snippet execution.

Monitoring Policies

- Updated the instructional language on the **Log File Monitoring Policies** page (System > Manage > Log File Monitoring Policies) to better guide the user.

Node Configuration Utility

- Updated the node configuration utility that you can access on an SL1 collector at `<IP address>:7700/node-config` using the em7admin account. This utility is used to configure the SL1 Collector and mimics the functionality of the web configurator.

- The node configuration utility includes a menu with four options: Home, Settings, Interfaces, and Connection.
 - The Home page displays information such as the type and IP address of the SL1 appliance.
 - The Settings page enables you to add and view proxy server information for SL1 appliances.
 - The Interfaces page displays an inventory of a node's interfaces. By clicking on an interface name, you can open the Interface Properties modal, which lets you view an interface's properties; update the interface name, IP address, gateway IP address, and netmask IP address; and create bonded interfaces. When you create bonded interfaces, the IP fields in the Bonding Interface modal will be validated to ensure that only valid IP formats are used; hostnames are not allowed.
 - The Connection page displays the SL1 appliance's connection status.
- You can now update the web configurator password by using the node configuration utility. To do this, log in to the node configuration utility, select the user name in the top-right corner of page, and click Change Password.

Nodes

- When you update the database password for an SL1 Collector, the password is encrypted.
- The **Registered Inventory** page (Manage > Nodes > Registered) was updated to indicate when a node is registered to multiple collector groups or organizations.
- Improved GraphQL error reporting for failed node token creation.
- When adding a node, you can add an identifying label when manually entering database address(es) to initiate a User Accepted Connection Request.
- If you attempt to create a System Accepted Connection in SL1 with an expired token, you are prompted to convert the expired token into a pending request.
- You can copy a Pending Request Confirmation Code to your clipboard. This six-character confirmation code is provided only after you either supply a VERIFY type token or enter a Database address manually, resulting in a phone home request.

Reports

- Updated the *Video Reports* PowerPack to ensure that all reports generate correctly when aligned with Dynamic Applications.

Run Book Actions and Automations

- Added a new **Execution Environment** drop-down field in the **Run Book Action Editor** for snippets to specify which Python version the Run Book Action should use. (If you currently have any existing Run Book Actions that use custom Python snippet code, see an important note in the section on [Libraries and Execution Environments](#).)
- Updated the **Execution Environment** drop-down for Custom Action Types to contain the concatenated Python version for the selected environment.

Security

- Updated packages to address the following "high" security vulnerabilities: ELSA-2022-6834, ELSA-2022-6878, ELSA-2022-9227, ELSA-2022-9421, ELSA-2022-9589.

Setup and Config Page

- Implemented a new **Setup and Config** page that provides guided journeys to help you configure and manage SL1. This page consists of the following options:
 - Get Started, which displays helpful onboarding information.
 - Overview, which allows you to begin the Setup & Configuration workflow journey. In this release, you have access to two workflow journeys: "Take a Tour of SL1" and "Discover and Monitor Hybrid Cloud Infrastructure".
 - Next Steps, where you can easily navigate to pages where you can manage devices, collector groups, organizations, users, and access hooks.
 - Resources, which provides links to helpful product support.
- The **Setup and Config** page automatically keeps track of your workflow progress. It also includes a **[Status]** button that allows you to proactively update the status of activity workflows. Statuses include "Not Started," "In Progress," "Complete," and "Not Applicable."
- You can access the **Setup and Configuration** page from the left navigation or the Advanced Menu, or by going to `<SL1 IP address>/setup-config`. To view this page, you must have the `SETUP_CONFIG_PAGE` access hook assigned to your user account and the `_GQL_SETUP_CONFIG` and `_AP2_SETUP_CONFIG` toggles enabled.

Subscription Billing

- Added the Legend Device Payload to the Subscription Usage Process.
- Improved the query performance in the Subscription Usage Crunch process.

System Administration

- For new installations of SL1 version 11.3.0 from the SL1 ISO file, the default MariaDB user will be "clientdbuser". When installing from the SL1 ISO file, the password entered in the install wizard for em7admin and root will also be used as the password for "clientdbuser" in MariaDB, but you can change the password during the initial installation. The "clientdbuser" user does not have root/superuser privileges.

NOTE: When deploying 11.3.0 on AWS using an AML cloud image, you must take additional manual steps to set up the "clientdbuser" password in MariaDB. For more information, see the [Known Issues](#) section.

- Adjusted file system permissions for the core application to remove "world read" permissions from most files.

- Enabled Military Unique Deployment (MUD) SL1 systems to use SSH key authentication without locking if the password expires.
- Added a local DNS cache to SL1 appliances to improve performance.

System Update

- A new post-update script checks for plain text credentials inside the "master.system_settings_licenses" table and encrypts the credentials if they are in plain text.
- Updated the System Status script, system_status.sh, for 11.3.0. This script provides diagnostic data for each appliance in your SL1 system.

User Accounts

- Added a new **User Display Name** column on the **User Accounts** page (Registry > Accounts > User Accounts). This is the user's name as it appears throughout SL1 and in system logs, as determined by the user's authentication resource settings.

User Interface

- You can now set dashboard and Advanced Menu pages as your preferred landing page when you log in to SL1. To do so, go to the page you want to set as your landing page, click your username in the upper right corner, and select **[Set As Landing Page]**.
- The following inventory pages and tabs in the SL1 user interface were updated to use a new set of filters for the columns in the list. You can start typing filter text or select filter options in one or more of these filters to narrow down the list to just the items you want to view:
 - Collections
 - Collector Groups
 - Credentials
 - Discovery Sessions
 - Classic Discovery
 - Machine Learning
 - Maps
 - Policies
 - Services
- Updated the **Collections** page in device details to include a component that enables you to see, search, filter, and select device collections.
- Updated the appearance of the **Severity** values that appear in tables for devices, events, and services.
- Updated the appearance and functionality of tables throughout SL1.

Webhooks

- You can now create a webhook receiver from a device's **Monitoring Policies** page. To do so, go to the **Monitoring Policies** page (Devices > Device Manager > wrench icon > Monitors) and click Create > Create Webhook Receiver. You can also edit and delete webhooks from this page.
- When you configure SL1 to receive webhook messages from third-party systems, you can now use a script to configure SSL certificates, specify ports, and enable nginx and firewall settings for webhooks.
- The **Create New Webhook** modal (Registry > Monitors > Webhooks > Create) now includes tooltips with brief descriptions of each field.

Issues Addressed in 11.3.0

This section describes the issues that were addressed in SL1 version 11.3.1.

API

- Addressed an issue in which invalid XML logic in the API Response class returned blank API key values. (Case: 00242634) (Case: 00271627) (Jira ID: EM-51073)
- Addressed an issue in which the /api/monitor process was not found after upgrading to SL1 11.2.0 (Case: 00273881) (Jira ID: EM-52808)

Backups

- Addressed an issue that was causing backups to fail when you scheduled the backup on an interval, such as daily or weekly. (Case: 00303104) (Jira ID: EM-54845)

Business Services

- Addressed an issue in which users could not sort or return results in Business Services when sorting by the **Category**, **Class**, or **Sub-Class** columns. (Case: 00235401) (Jira ID: EM-49253)
- Addressed an issue in which the 'Save_topology' script returned incorrect constituents when the Device Service filter read as a dynamic Device Group with a dynamic rule. (Case: 00219151) (Jira ID: SLUI-10883)
- Addressed an issue for which only the last added device appeared in the **Selected Devices** field of the Service Usage Policy editor. All available devices now appear as expected. (Case: 00214808) (Jira ID: EM-47809)
- Fixed a regression that prevented the URL parameter 'harProviderId' from being passed to its linked service dashboards; this was causing the Service View widget to load incorrectly. (Case: 00218610) (Jira ID: EM-48269)
- Addressed an issue that was causing unhandled exception errors for the IT Service: Service Management Engine process. (Case: 00235152) (Case: 00264093) (Jira ID: EM-49474)

Credentials

- Addressed an issue in which SOAP/XML credentials were not translating the variable %N as the device hostname. (Case: 00219275) (Jira ID: EM-48235)
- Updated PowerShell credentials to ensure the credentials are accepted by the server and can collect data properly. (Case: 00257763)(Jira ID: EM-52647)

Custom Attributes

- Updated custom attribute mapping so that custom attributes are no longer associated with a device after the device has been deleted. (Case: 00226938) (Jira ID: EM-48682)
- Addressed an issue in which device custom attributes failed to display in the Events table on the **Dashboards** page. (Case: 00240781) (Jira ID: EM-46922)

Daily Maintenance

- Updated the Daily Maintenance database connection to eliminate timeouts due to inactivity. (Case: 00236344) (Jira ID: EM-51173)

Dashboards

- Addressed an issue in which the Dashboard Journal data table failed to display in the classic SL1 user interface. (Case: 00007561) (Jira ID: EM-32948)
- Scheduled dashboards in the classic SL1 user interface were updated to ensure that they export and email to users as intended, and their margin widths were adjusted to ensure that they display properly. (Case: 00133754) (Case: 00251107) (Jira ID: EM-51659) (Jira ID: EM-40725)

Data Collection

- Non-active Database Server and Data Engine appliances now periodically check collector connectivity. If a collector is not reachable, SL1 raises an event listing the collector and which Database Server or Data Engine cannot reach it. (Case: 00096310) (Jira ID: EM-37453)
- On SL1 Appliances, Docker network bridge mode is disabled by default. This setting lets SL1 monitor devices other than internal Docker components that might be on the 172.17.0.0/16 network. (Cases: 00208448, 00251943, 00273269) (Jira ID: EM-47339)

Data Retention

- Updated the following Template retention settings to allow a minimal value setting of 1:
 - Hourly Rollup Performance Data on the Config tab
 - Hourly Rollup Bandwidth Data on the Config tab
 - Hourly Rollup Retention on the Dynamic Applications tab (Case:00153850) (Jira ID: EM-41927)
- Updated the following Device Editor retention settings to allow a minimal value setting of 1:

- Hourly Rollup Performance Data from the Threshold Tab in Data Retention Thresholds
- Hourly Rollup Bandwidth Data from the Threshold Tab in Data Retention Thresholds
- Hourly Rollup Retention from the Threshold Tab in any Dynamic Application where this is defined (Case:00153850) (Jira ID: EM-41927)
- Updated the PruneInventoryData task to include an hourly data pruner that deletes expired master_dev.device_services data to address an issue in which "Svc Policies" did not load. (Case: 00187698) (Jira ID: EM-47343)
- Updated the PruneInventoryData task to include an hourly data pruner to delete expired master_dev.device_ports. (Case: 00028368) (Jira ID: EM-31142)
- Updated the Hourly Maintenance Device Pruner to be more efficient and purge a list of devices instead of a single device. Purging the devices tasks is set as the final task to execute, so other hourly maintenance tasks are finished first. (Case: 00256727) (Jira ID: EM-51415)

Devices

- Addressed an issue in which the Device Reports panel was displaying the incorrect status for Windows services. (Case: 00039429) (Jira ID: EM-32376)
- Improved device deletion performance. (Case: 00268318) (Jira ID: EM-52467)
- Updated the Device Investigator to ensure that it returns accurate Dynamic Application data. (Case: 00236115) (JIRA ID: EM-13554)
- Addressed an issue in the Device REST API that caused an erroneous "User Maintenance enabled for device via API" message in the device log. (Case: 00235156) (Jira ID: EM-49614)
- Addressed an issue where user-initiated Run Book Automations were not displaying in the Tools widget for event-based device summaries. (Case: 00279219) (Jira ID: EM-52991)

Device Templates

- Updated the **Device Template Editor** so you can apply the template only to a device group that has matched devices. You cannot apply a template to an empty device group. (Case: 00088025) (Case: 00098339) (Jira ID: EM-25170)

Documentation

- Updated the **Events** manual to provide additional details about how SL1 determines the yName value in an event message based on the event's source type. (Case: 00140767) (Jira ID: EM-42421)

Dynamic Applications

- Added an index to a Dynamic Application alerts database table to improve query performance. (Case: 00258275) (Jira ID: EM-51388) (Jira ID: EM-51520)
- Addressed an issue in which hourly maintenance was removing Dynamic Application alignments on vanished devices. (Case: 00218839) (Jira ID: EM-49258)

- Made enhancements to collection object deletion and data removal using the **Dynamic Application Collections** page (Devices > Device > Device Manager > wrench icon > Collections) to prevent performance issues when deleting a large amount of collected data. (Case: 00159668) (Jira ID: EM-42726)
- Made improvements to the alert function for standard deviation (deviation). Added concurrent threads to the deviation crunch process (da_deviation_crunch.py) and added an index to the database table dynamic_app_data_x.dev_stats_y. (Case: 00234683) (Jira ID: EM-49204) (Jira ID: EM-49799) (Jira ID: EM-50615) (Jira ID: EM-50616)
- Updated the "Linux: CPU Performance" Dynamic Application to address an issue that was causing the performance view report and device utilization report to display different CPU utilization values for the same Linux server devices. (Case: 00233235) (Jira ID: EM-48304) (Jira ID: EM-49429)
- Improved the query for collection objects in the **Dynamic Application Collections** page (Registry > Devices > wrench icon > Collections) to prevent timeouts. (Case: 00199764) (Case: 00207217) (Case: 00209781) (Jira ID: EM-47087) (Jira ID: EM-46937) (Jira ID: EM-46193)

Events

- Updated the Events API so the "isnull" filter recognizes values of 0 in the user_ack field. (Case: 00208655) (Jira ID: EM-47576)
- Addressed an issue in which event processing failed due to the Event Engine restarting every few minutes. The default memory limit for the Event Engine has been increased from 1 GB to 2 GB. (Case: 00238925) (Jira ID: EM-43969)
- Updated the search and filtering functionality for the **Auto Clear** field on the **[Advanced]** tab of the **Event Policy Editor** to ensure that returns the correct events list. (Case: 00282405) (Jira ID: 53247)

Form Fields

- Updated the Custom Form Fields to be editable and to allow special characters in the name of the access key. (Case: 00250454) (Jira ID: EM-50846)

GraphQL

- The GraphQL query for deviceRelationships now returns the correct DCM+R relationship type. (Case: 00227848) (Jira ID: EM-50353) (Jira ID: EM-49546) (Jira ID: SLUI-14078) (Jira ID: EM-49546) (Jira ID: EM-49503)
- Addressed an issue where GraphQL responses contained extremely large cursor sizes. SL1 no longer appends search objects to the cursor. (Case: 00240115) (Jira ID: EM-49723) (Jira ID: EM-50126) (Jira ID: 50193) (Jira ID: EM-49316)
- Updated the GraphQL custom attribute search on devices to be performant with Aurora's query. (Case: 00232092) (Jira ID: EM-49225)
- Addressed an issue where searching device classes on devices sometimes showed invalid results. (Case: 00233392) (Jira ID: EM-49180)
- Updated the GraphQL back-end so that the legend_device table alias 'ld' registers with the framework. It is now safely referenced in the device_state subquery. (Case: 00223222) (Jira ID: EM-50746)

Interface Billing

- Improved organization restrictions on user visibility for the widgets on the Interface Billing dashboard to ensure widgets display the correct information relevant to the Interface Billing Policy selected in the Context Driver widget. (Case: 00207171)(Jira ID: SOL-18166)

Monitoring Policies

- TCP Port Monitor policies will no longer return unhandled exceptions for policy key errors. (Case: 00238633) (Jira ID: EM-51305)

Platform

- Updated the Enterprise Database: Collector Config Push process (config_push.py) to improve table view cleanup. (Case: 00228062) (Jira ID: EM-48907)
- Updated the name and default value of the setting used to define the wait time for Config Push result messages. The name was changed from "message_timeout" to "result_wait_timeout". This update only affects users who have manually added the "message_timeout" configuration to a "[CONFIG_PUSH]" section of their "/etc/silo.conf" configuration file on their central database. Users need to update the timeout name. The default timeout value was changed from "60" to "300". (Case: 00283693) (Jira ID: EM-53433)
- Added a configurable dcm_get_lock_timeout value to the master.system_custom_config database table that can be updated to change the timeout for the GET_LOCK query that is called before creating and updating device components. If the GET_LOCK query is causing significant load contention on your SL1 system, contact ScienceLogic Support to adjust this new database value. (Case: 00257030) (Jira ID: EM-51271) (Jira ID: EM-51642)
- Enhanced the performance of the UPDATE queries against the master.dynamic_app_collection database table. (Case: 00257038) (Jira ID: EM-51272) (Jira ID: EM-51367)
- Updated the Config Push lookup table schema to support collector group ID values greater than 255. (Case: 00295890) (Jira ID: EM-54213)

PowerShell Collection

- Updated the Docker image version that is included with SL1 to address an issue that was causing concurrent PowerShell collections to fail on some Data Collectors that had more than 400 devices. (Jira ID: EM-54212)

Reports

- Updated the **Archived Job** page (Reports > Create Report > Scheduled Job/Report Archive > Archived Job) to display the correct count of archived jobs at the top of the page. (Case: 00225414) (Jira ID: EM-48859)
- Updated the **Ad-hoc and Scheduled Reports** page (System > Settings > Data Retention > Ad-hoc and Scheduled Reports) to return accurate data during a query. (Case: 00225414) (Jira ID: EM-48859)

Run Book Actions and Automations

- Added Cross-Site Request Forgery (CSRF) Validation to the **Run Book Action Policy Manager** page (Registry > Run Book > Actions) for improved security. (Case: 00231093) (Jira ID: EM-49044)

Schedules

- Updated the Automation Schedule to set the owner to "em7admin" upon PowerPack installation. (Case: 00245771) (Jira ID: EM-50620)
- Updated schedules in SL1 to properly identify non-recurring, expired schedules and ensure those expired schedules are removed. (Case: 00236585) (Jira ID: EM-48629)

Subscription Billing

- Updated the Feature Usage Data to ensure it successfully connects to a proxy system via HTTP or HTTPS. (Case: 00256955) (Jira ID: EM-51786)

System Administration

- Updated the visilo file to be able to handle /etc/silo.conf file with unencrypted passwords. (Case: 00286881) (Jira ID: EM-53481)

User Accounts

- Addressed an issue in which API restrictions were not enforced for different user accounts. User accounts with different access permissions shared the same API-access despite their accounts' assigned access permissions. (Case: 00146238) (Jira ID: EM-42811)

User Interface

- Addressed an issue where the platform title as selected in the user's selected theme was not displaying correctly in the user's browser. (Case: 00235441) (Jira ID: EM-49280)
- Added new access hooks to control visibility over pages. With this update, only administrators can see the **Administer Bookmarks** (Misc > Bookmarks), **Regular Expression Tester** (Misc > Regex Tester), and **SL1 License Info** (Misc > License Info) pages. (Case: 00253683) (Jira ID: EM-51840)

Installing and Upgrading SL1

For a detailed overview of SL1, see the [Introduction to SL1](#) manual.

For detailed instructions on performing a new installation of SL1, see the [Installation and Initial Configuration](#) manual.

For detailed instructions on upgrading SL1, see the chapter on "Updating SL1" in the [System Administration](#) manual and the upgrade notes that are included in this document.

NOTE: ScienceLogic strongly recommends that you review the [Known Issues](https://support.sciencelogic.com/s/known-issues#sort=relevancy) for SL1 (<https://support.sciencelogic.com/s/known-issues#sort=relevancy>) before installing a new update. For known issues specific to this release, see the [Known Issues](#) section of this document.

Important Upgrade Notes

This section includes important notes for upgrading existing SL1 systems to the 11.3.1 release.

Unless otherwise stated, the information in this section applies to all users who are upgrading from previous SL1 versions.

CAUTION: ScienceLogic strongly recommends that you review these upgrade notes in their entirety before upgrading to version 11.3.1.

Supported Upgrade Paths

Previous SL1 releases included major updates that you must consume before you can upgrade to 11.x. Therefore, you might be required to upgrade to one or more earlier versions of SL1 before you can upgrade to 11.x.

Depending on the version of SL1 that you are currently running, you can upgrade to SL1 11.x using one of the following upgrade paths:

- 11.x to 11.x
- 10.x to 11.x
- 8.12 to 10.x to 11.x
- 8.10 to 8.12 to 10.x to 11.x
- 8.9 to 8.10 to 8.12 to 10.x to 11.x

Be advised that *upgrading from a few specific versions is not supported*.

If you are currently running a version prior to 10.1.0, then you must also upgrade to the [MariaDB version that corresponds to each SL1 release in your upgrade path](#).

If you are currently running SL1 version 10.1.0 or later, you can upgrade directly to SL1 version 11.1.x and the corresponding MariaDB version.

Before upgrading between SL1 versions, contact ScienceLogic Support to ensure that the upgrade paths between those versions is supported.

Unsupported Upgrade Paths

If your SL1 system is deployed on AWS, you cannot upgrade from the following versions to SL1 11.3.1 at this time due to a known issue:

- 10.2.5
- 10.2.6
- 10.2.6.1
- 10.2.7
- 11.1.3
- 11.1.4
- 11.1.5
- 11.1.6

This known issue does not impact deployment types other than AWS. A resolution is planned for an upcoming release.

Upgrading MariaDB and Rebooting SL1

CAUTION: A known issue might cause the MariaDB upgrade script to fail if you are updating from an SL1 11.3.0 ISO to 11.3.1. For more information and a workaround, see the section [Upgrading MariaDB for SL1 11.3.0 to 11.3.1 Upgrades](#).

CAUTION: Due to the upgrade to MariaDB 10.4.28 that is required to consume SL1 11.3.1, you can no longer execute multi-statement SQL operations in Python in this release. This issue should only impact snippets and Dynamic Applications that use custom code with statements like `"dbc.execute('SELECT 1; SELECT 2; SELECT 3')"`, which was a practice that was already discouraged due to the excess load such statements can cause on the Database Server.

Some SL1 versions include important security updates. To apply these updates, you must upgrade MariaDB and then reboot all SL1 appliances.

The following table specifies the required MariaDB version for each SL1 version and which SL1 updates require you to reboot all SL1 appliances:

SL1 Release	Required MariaDB Version	Requires Appliance Reboot?
11.3.1	10.4.28	Yes
11.3.0	10.4.26	Yes
11.2.3	10.4.28	Yes
11.2.2	10.4.26	Yes

SL1 Release	Required MariaDB Version	Requires Appliance Reboot?
11.2.0	10.4.24	Yes
11.1.6	10.4.28	Yes
11.1.5	10.4.26	Yes
11.1.4	10.4.26	Yes
11.1.3	10.4.25	Yes
11.1.2	10.4.24	Yes
11.1.1	10.4.22	Yes
11.1.0	10.4.20	Yes
10.2.7	10.4.27	Yes
10.2.6.1	10.4.26	Yes
10.2.6	10.4.26	Yes
10.2.5	10.4.22	Yes
10.2.4.1	10.4.22	Yes
10.2.4	10.4.22	Yes
10.2.3	10.4.21	Yes
10.2.2	10.4.18	Yes
10.2.1	10.4.18	Yes
10.2.0	10.4.18	Yes

NOTE: Your MariaDB password can have a maximum of 23 characters due to a [known issue](#).

NOTE: For instructions on updating MariaDB or rebooting the SL1 system, see the section on [Updating SL1](#) in the [System Administration](#) manual.

If you would like assistance in planning an upgrade path that meets your security needs while minimizing downtime, please contact your Customer Success Manager.

Upgrading MariaDB for SL1 11.3.0 to 11.3.1 Upgrades

A known issue might cause the MariaDB upgrade script to fail if you are updating from an SL1 11.3.0 ISO to 11.3.1.

To work around this issue, prior to running the `module_upgrade_mariadb` script, enter the mysql prompt as "root" and run the following command:

```
GRANT ALL PRIVILEGES ON *.* TO 'root'@'localhost' IDENTIFIED VIA unix_socket WITH GRANT OPTION;
```

You can then run the `module_upgrade_mariadb` script as normal.

System Update Notes

- **SL1 updates overwrite changes to the configuration file `/opt/em7/nextui/nextui.env`.** This is a known issue. (For more details, see <https://support.sciencelogic.com/s/article/1161> and <https://support.sciencelogic.com/s/article/1423>.) ScienceLogic recommends that you back up this file before applying an update and then reapply your changes to this file.
- During the normal system update process, multiple processes are stopped and restarted. This might result in missed polls, gaps in data, and/or unexpected errors. ScienceLogic recommends that you always install SL1 releases during a maintenance window.
- The SL1 system update process starts a background process that can perform any required post-upgrade tasks. The post-patch update process is automatically stopped after 24 hours. However, depending on the size of your database as well as the version from which you are upgrading, the post-upgrade tasks can take several days to perform. If the post-patch update process is stopped after 24 hours, the process will automatically re-start and continue processing from the point at which it was stopped. If you see an event that indicates the post-patch update process was stopped, you do not need to contact ScienceLogic support for assistance until you see the same event for three consecutive days.

Required PowerPack Updates

Required Version Updates

If you are using the following PowerPacks, you must upgrade to the specified minimum supported versions before upgrading to SL1 version 11.3.x:

- Cisco: ACI v112
- Cisco: AppDynamics v102
- Cisco: Cloud Services Platform v107
- Cisco: Viptela v104
- Datacenter Advanced Enrichment Actions v106
- Dynatrace v105
- HTTP Action Type v103
- IBM: DB2 v104
- Kubernetes v104
- Linux: Base Pack v105
- Linux SSH Automation v104
- Microsoft: Azure v115
- Microsoft: Office 365 v106
- NetApp: Base Pack v106
- Oracle: MySQL v102

- VMware Automation v102
- Windows PowerShell Automation v104

Earlier versions of these PowerPacks will not prevent SL1 versions 11.3.x from installing or operating, but they might not operate as expected after the SL1 upgrade due to technical incompatibilities.

Required Credential Updates

Some PowerPacks require you to update their credentials before you upgrade to version 11.2.0 or later. Therefore, if you are using one of the following PowerPacks, you must edit an HTTP header in the credential before you upgrade to version 11.3.x:

- Cisco: ACI Multisite
- CouchBase
- Dell: EMC VMAX
- Google: Cloud Platform
- LayerX: Appliance Monitoring
- ScienceLogic: PowerFlow
- PowerPacks built using the REST PowerPack

To edit the credential HTTP header:

1. Go to the **Credentials** page (Manage > Credentials).
2. Locate the credential you created, then click its **[Actions]** icon and select *Edit/Test*.
3. Find the "Content-Type: application/json" HTTP header, then remove the space in the HTTP header so that the new header reads "Content-Type:application/json".
4. Repeat step 3 for any other HTTP header entries in the credential.
5. Click **[Save & Close]**.
6. Repeat these steps for any other credential relating to the PowerPacks in the list above.

Required Updates for Users Running Amazon RDS (Aurora MySQL 5.7)

If you are using Amazon RDS (Aurora MySQL 5.7) with SL1 and are upgrading from a version of SL1 prior to 11.2.0, then you must update to the following PowerPack versions before installing SL1 version 11.3.x:

- Cisco: UC VOS Applications v110

If you are using Amazon RDS (Aurora MySQL 5.6) with SL1, older versions of this PowerPack will continue to work with SL1 version 11.3.x.

New and Updated PowerPacks

The 11.3.1 ISO includes the following new or updated PowerPacks:

- Interface Billing v101
- Linux Base Pack v107
- Microsoft: Azure v116
- Microsoft Base Pack v107

- Microsoft: Windows Server v114
- Video Reports v103

Before consuming SL1 11.3.1, please verify whether any PowerPacks currently running on your system are newer than the PowerPacks included in this release. If the PowerPack on your system is newer than the one included with this release, you might see spurious error messages. To avoid spurious error messages:

1. Go to the **Device Components** page (Registry > Devices > Device Components).
2. Find each root device associated with the PowerPacks you do not want to update and select their checkboxes.
3. Click the **Select Action** field and choose *Change Collection State: Disabled (recursive)*, and then click the **[Go]** button.
4. Wait five minutes after disabling collection.
5. Install the SL1 update.
6. Go to the **Device Components** page (Registry > Devices > Device Components).
7. Select the checkbox for all affected root devices.
8. Click the **Select Action** field and choose *Change Collection State: Enabled (recursive)*, and then click the **[Go]** button.

PowerPacks Removed from the ISO

NOTE: If you are upgrading from a previous version of SL1, the 11.3.1 upgrade will not remove any existing PowerPacks. The PowerPacks listed below are still available for download from the [PowerPacks Support](#) page.

The 11.3.1 release removes the following PowerPacks from the ISO:

- Acano MCU
- Cisco: ACI
- Cisco: Cloud Services Platform
- Cisco: Contact Center Enterprise
- Cisco: CUCM
- Cisco: Hyperflex
- Cisco: Medianet/Mediatrace
- Cisco: TelePresence Conductor
- Cisco: TelePresence: Endpoints
- Cisco: TelePresence: Infrastructure
- Cisco: UC Ancillary
- Cisco: UC VOS Applications
- Cisco: UCS

- Cisco: UCS Director
- Cisco: Unity
- Cisco: Video Endpoint
- Entity MIB
- LayerX Appliance Monitoring
- Microsoft: IIS Server
- Microsoft: Lync Server 2013
- Microsoft: Lync Server 2013 Dashboards
- Polycom Endpoint
- Polycom: Infrastructure
- SL1: Concurrent PowerShell Monitor
- Tandberg: Infrastructure

Monitoring Windows with WMI

NOTE: This section applies to users who are upgrading from the following releases:

- SL1 11.1.0 through 11.1.2
- Any SL1 release prior to 10.2.5

If you are upgrading from the following releases, you can ignore this section:

- 11.2.0 and later
- 11.1.3 or a later 11.1.x version
- 10.2.5 or a later 10.2.x version

SL1 versions 11.2.0, 11.1.3, and 10.2.5 included a new WMI client in response to Microsoft security updates. This change enables WMI Dynamic Applications to collect data from hardened Windows servers, but also has a major impact on system scalability.

This change significantly decreases the number of Microsoft Windows servers that can be supported on each Data Collector in your SL1 system. Users who need to monitor Windows devices using WMI should analyze their system resources and capacity before upgrading to 11.3.1. For guidance about sizing, see the updated [Collector Sizing guidelines for WMI endpoints](#).

To avoid this impact, ScienceLogic recommends using SNMP collection for two-core Windows servers and PowerShell collection for four-core Windows servers. For more information, see this [Support Knowledge Base article](#).

Upgrading to 11.3.1 with Amazon RDS

NOTE: This section applies to users who are upgrading from SL1 11.1.x or earlier.

If you are upgrading from SL1 11.2.0 or later, you can ignore this section.

A known issue is causing an incompatibility between the SL1 user interface in versions 11.2.x and later and versions of the Amazon RDS (Aurora MySQL) database that were supported by SL1 versions 11.1.x and earlier.

If your SL1 system uses Amazon RDS for storage and you plan to upgrade to SL1 11.3.1 from a version prior to 11.2.0, then immediately after upgrading SL1, you must also upgrade the MySQL database to version 5.7 with an Aurora MySQL engine of 2.x.

If you upgrade to SL1 11.3.1 but do not immediately upgrade the MySQL database to version 5.7 with an Aurora MySQL engine of 2.x, then you will experience an outage and the SL1 user interface will be unavailable until MySQL is upgraded.

To plan an upgrade strategy that minimizes downtime, ScienceLogic strongly suggests that you contact your Customer Success Manager before upgrading to SL1 11.3.1 and upgrading Amazon RDS (Aurora MySQL).

Pre-Upgrade Test for PhoneHome Database Servers

NOTE: This section applies to users who are upgrading from SL1 11.1.x or earlier and have an existing PhoneHome configuration.

If you are upgrading from SL1 11.2.0 or later or you do not have a pre-11.2.0 PhoneHome configuration, you can ignore this section.

SL1 version 11.2.0 included a new pre-upgrade test that checks for existing PhoneHome Database Servers.

This pre-upgrade test looks for PhoneHome token IDs inside the `/home/phonehome0/config.json` file and fails if the value of the token ID field is less than or equal to "0". In previous versions of SL1, the primary PhoneHome Database was not self-registered with a token, causing it to have an ID of "0".

Therefore, if you are upgrading from version 11.1.x or earlier and you have a PhoneHome configuration, then you must perform these one-time manual configuration steps on all Database Servers in your PhoneHome configuration prior to upgrading to SL1 version 11.3.1:

1. Log in to the console of the Database Server or use SSH to access the server.
2. To determine if all of your PhoneHome Database Servers are registered, type the following command and check if any have an ID value of "0":

```
cat /home/phonehome0/config.json
```

3. If a PhoneHome Database Server has an ID value of "0", type the following command and locate the ID of the current appliance:

```
phonehome status
```

4. Type the following command and locate the PhoneHome token:

```
phonehome token <ID from step 3>
```

5. Type the following command to register the PhoneHome token:

```
phonehome register <token from step 4>
```

6. Repeat steps 3-5 for all PhoneHome Database Servers that have an ID value of "0".

7. Type the following command to ensure that all of your PhoneHome Database Servers are synced:

```
phonehome sync
```

8. Repeat step 2 and confirm that all Database Servers have ID values greater than "0".

NOTE: Do not attempt to upgrade to 11.3.1 until all pre-upgrade tests are successful on all PhoneHome Database Servers.

IMPORTANT: The PhoneHome server process runs as an unprivileged user that will not be able to bind to a privileged port (1-1023). Therefore, when you choose a custom port, you must choose port 1024 or higher.

PHP Updates

NOTE: This section applies to users who are upgrading from SL1 10.2.x or earlier.

If you are upgrading from SL1 11.1.0 or later, you can ignore this section.

In SL1 version 11.1.0, all PHP code was converted to PHP 7. Therefore, if you are upgrading from a version of SL1 prior to 11.1.0, please note the following:

- During the upgrade to 11.3.1, the user interface will be unavailable for several minutes.
- Versions of Global Manager prior to 11.1.0 will not work with SL1 11.1.0 or later.
- Web Proxy Services will not work in SL1 11.1.0 or later.
- PowerPacks built in SL1 version 11.1.0 and later releases cannot be imported into previous versions of SL1. However, PowerPacks built in releases prior to 11.1.0 can be imported into 11.1.0 and later.
- If you have created custom content in PHP, see this page for notes on backward compatibility:
<https://www.php.net/manual/en/migration70.incompatible.php>

Required Updates When Upgrading from Version 8.14.x or Earlier

NOTE: This section applies to users who are upgrading from SL1 8.14.x or earlier.

If you are upgrading from SL1 10.1.0 or later, you can ignore this section.

To avoid a known issue, if you are currently running SL1 8.14.x or earlier, you **must** first upgrade to the latest SL1 10.2.x release and the version of MariaDB that corresponds to that release before upgrading to SL1 version 11.3.1 and MariaDB version 10.4.28.

Therefore, if you are currently running version 8.14.x or earlier, then the upgrade to SL1 version 11.3.1 might require up to four (4) maintenance windows:

1. Upgrade to SL1 version 10.2.x or later.
2. Upgrade to the [MariaDB version that corresponds to that SL1 release](#).
3. Upgrade to SL1 version 11.3.1.
4. Upgrade to MariaDB version 10.4.28.

For detailed instructions on planning an upgrade, best practices for upgrades, and executing an upgrade, see the section on [Updating SL1](#) in the [System Administration](#) manual.

Other Considerations When Upgrading from Version 8.14.x or Earlier

NOTE: This section applies to users who are upgrading from SL1 8.14.x or earlier.

If you are upgrading from SL1 10.1.0 or later, you can ignore this section.

Consider the following additional notes before deploying 11.3.1:

- As of version 10.1.0, SL1 no longer includes Flash.
- As of SL1 8.12.2, ScienceLogic no longer updates the help that appears when you click the **[Guide]** button that appears in the classic user interface. The SL1 user interface provides a new tool for inline help. Clicking the **[Help]** button at the top of the page now opens a Help topic about that page on the right-hand side of the page. From that topic, you can also click a context-sensitive link to open the relevant page in the product documentation at docs.sciencelogic.com in a new browser window.
- As of SL1 8.10.0, SL1 does not support Data Collectors and Message Collectors running the CentOS operating system. If your system includes Data Collectors and Message Collectors running the CentOS operating system, contact your Customer Success Manager for details on upgrading Data Collectors and Message Collectors to Oracle Linux before installing the latest SL1 version.
- To download updates for previous SL1 versions that have reached their End-of-Life date and are no longer supported by ScienceLogic, contact ScienceLogic Support or a designated Customer Success Manager to get the update files.

Known Issues for SL1 Version 11.3.1

NOTE: ScienceLogic strongly recommends that you review all [Known Issues](#) for SL1. For more information, see <https://support.sciencelogic.com/s/known-issues#sort=relevancy>.

The following known issues exist for SL1 version 11.3.1:

- Due to a known issue in AWS GovCloud deployments in Military Unique Deployment (MUD) systems, some automation engine services, such as Run Book automations, may not be available during an HA failover. (Jira ID: EM-54991)
- A known issue might cause the MariaDB upgrade script to fail if you are updating from an SL1 11.3.0 ISO to 11.3.1. For more information and a workaround, see the section [Upgrading MariaDB for SL1 11.3.0 to 11.3.1 Upgrades](#). (Jira ID: EM-57702)
- MariaDB passwords are currently limited to 23 characters due to an issue that truncates longer passwords in the master.system_settings_licenses database table. If your MariaDB password is longer than 23 characters, your SL1 appliances will not update properly and will appear to be unavailable. To work around this issue, shorten your MariaDB password to 23 characters or less. (Jira ID: EM-58684)
- Due to a known issue with the logrotate utility, there is a chance that the log files of processes with debug mode set to "enabled" will grow quickly without rotating. This can fill up the /var/log partition, causing issues with further logging and patching of the affected collector appliances. (Jira ID: EM-43992)
- When you filter the Agents page using the column heading filters at the top of the page, if you clear the values to see the complete list again, the table is not reloading the entire list. To work around this issue, you must refresh the page to clear the filter. (Jira ID: EM-57824)
- After an upgrade from an SL1 11.x release to version 11.3.1, you might experience an issue that prevents you from logging in to the user interface due to the system reporting that your SL1 appliance is not licensed. This is happening in two specific upgrade scenarios:
 - New customers who have never had a permanent license and are upgrading from the 11.2.0 ISO to 11.3.1
 - Existing customers who previously upgraded from an earlier release to 11.2.0 or 11.2.1 and are now upgrading to 11.3.1 and whose original license expiry date is prior to the date they upgrade to 11.3.1

For additional details about this issue and the workaround, see:

<https://support.sciencelogic.com/s/article/9617>. (Jira ID: EM-54958)

- When deploying SL1 on AWS using the 11.3.1 AML, you must take additional manual steps to set up the "clientdbuser" password in MariaDB. This requires you to edit the /etc/silo.conf file. If you do not, you will not be able to access the user interface and a banner message about the database password not being set will appear after you log in to the appliance using SSH. For additional details about this issue, see: <https://support.sciencelogic.com/s/article/9875>.
- When executing AWS guided workflow discoveries, executing the same workflow or workflows with similar settings can result in asset duplication. For IAM guided workflows, this will result in completely duplicated

account device component trees. For other AWS workflows, this might result in duplicated virtual devices that represent the AWS organization. (Jira ID: EM-54284)

- You cannot delete a PowerPack that has a PowerShell credential. To work around this issue, delete all PowerShell credentials from that PowerPack, and then delete the PowerPack. (Jira ID: EM-53453)
- If you repeatedly sign in and out of SL1 in a short period of time, you might receive an error that temporarily prevents you from signing back in due to a caching issue. If this occurs, you can try one of the following workarounds:(Jira ID: SLUI-15357)
 - Wait 5 minutes before attempting to sign in again.
 - Set caching for SL1 sessions to 0. Doing so avoids the issue by effectively turning off session caching, but this might result in performance issues or issues with Global Manager.
 - Sign in using the classic SL1 user interface.

For details, see <https://support.sciencelogic.com/s/article/9715>.

- With on-premises SL1 Extended systems, the TLS handshake can fail between the Windows Agent on a monitored device and the SL1 streamer service. For details and the workaround, see: <https://support.sciencelogic.com/s/article/7417>. (Jira IDs: DO-4079, DO-4115)
- If you use CAC with LDAP/AD, then you must define an LDAP service account with permissions that allow the service account to query LDAP before configuring CAC. (Case: 00207174) (Jira ID: EM-46927)

SL1 Extended Architecture

NOTE: New installations of SL1 Extended Architecture are available only on SaaS deployments. For existing on-premises deployments of SL1 Extended Architecture, see the section on [Upgrading the SL1 Extended Architecture](#) in the [System Administration](#) manual.

11.3.1 supports the SL1 Extended Architecture. **The following SL1 features require the SL1 Extended Architecture:**

- **Anomaly Detection and future AI/ML developments.** Anomaly detection is a technique that uses machine learning to identify unusual patterns that do not conform to expected behavior. SL1 does this by collecting data for a particular metric over a period of time, learning the patterns of that particular device metric, and then choosing the best possible algorithm to analyze that data. Anomalies are detected when the actual collected data value falls outside the boundaries of the expected value range.
- **Expanded Agent Capabilities.** You can configure the SL1 Agent to communicate with SL1 via a dedicated Message Collector. However, this configuration limits the capabilities of the SL1 Agent. If you configure the SL1 Agent to communicate with SL1 via a Compute Cluster, you expand the capabilities of the SL1 Agent to include features like extensible collection and application monitoring.

- **Data Pipelines.** Data pipelines transport and transform data. Data transformations include enrichment with metadata, data rollup, and pattern-matching for alerting and automation. The Data Pipelines provide an alternative to the existing methods of data transport in SL1 (data pull, config push, streamer, and communication via encrypted SQL). Data pipelines introduce message queues and communicate using encrypted web services.
- **Publisher.** Publisher enables the egress of data from SL1. Publisher can provide data for long-term storage or provide input to other applications that perform analysis or reporting.
- **Scale-out storage of performance data.** Extended Architecture includes a non-SQL database (Scylla) for scalable storage of performance data.

The SL1 Extended Architecture includes four additional types of SL1 Appliances:

- **Compute Cluster.** Compute nodes are the SL1 appliances that run services that transport, process, and consume the data from Data Collectors and the SL1 Agent. SL1 uses Docker and Kubernetes to deploy and manage these services.
- **Load Balancer.** The SL1 appliance that brokers communication with services running on the Compute Cluster. Services running on the Compute Cluster are managed by Kubernetes. Therefore, a single service could be running on one Compute node in the Compute Cluster; to provide scale, multiple instances of a single service could be running on one, many, or all nodes in the Compute Cluster. To provide scale and resiliency, you can include multiple Load Balancers in your configuration.
- **Storage Cluster.** SL1 Extended includes a Storage Cluster that includes multiple Storage Nodes and one Storage Manager node. These SL1 appliances provide a NoSQL alternative to the SL1 relational database. The Storage Cluster can store performance and log data collected by the Data Collectors and the SL1 Agent.
- **Management Node.** The Management Node allows administrators to install, configure, and update packages on the Compute Nodes cluster, Storage Nodes, and the Load Balancer. The Management Node also allows administrators to deploy and update services running on the Compute Cluster.
- Resiliency and redundancy can also be accomplished by adding additional appliances to these configurations.

<p>NOTE: SL1 Extended Architecture does not provide MUD support.</p>

Recently Deprecated Features

11.3.0

- The 11.3.0 release deprecates the following PowerPack and removes it from the ISO:
 - SL1: Concurrent PowerShell Monitor v103

11.2.0

- The Inbox feature is deprecated and is available only for reports in the classic user interface.
- The IPMI features were deprecated in SL1.
- The Knowledge Base and Knowledge Base tab were removed from SL1.
- The clipboard feature was removed from Ticketing. The access hook for the clipboard feature was also removed.

11.1.0

- Removed Flash-based pages in Views > Other Views.
- The Flash-based System Usage Pie Chart has been deprecated and is no longer available on the system usage report.
- Removed Flash-based Hardware Inventory graph.
- Removed Flash-based Maps.
- Removed the Flash-based Org Clock.
- Removed the Flash-based Map Icon column in the Device Category Register.
- Removed the Flash-based Ticket Timeline report.
- Removed the Flash-based Event Overview report.
- The 11.1.0 release deprecated the following PowerPacks and removed them from the ISO:
 - Cisco: CUCM Dashboards
 - Cisco: Old Cisco Apps
 - Cisco Unity Pack
 - LayerX Cisco CDR
 - Link Layer Neighbor Discovery
 - Microsoft: Azure Classic
 - Microsoft: Exchange Server 2010
 - Microsoft: Exchange Server 2010 Dashboards
 - Microsoft: Lync Server 2010
 - Microsoft: Lync Server 2010 Dashboards
 - Microsoft: Windows Server Services. Its content now resides in the Microsoft Windows Server PowerPack v112 or later.

© 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/company/legal>.



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

International: +1-703-354-1010