



---

## SL1 11.1.0 Release Notes

SL1 version 11.1.0 (Document Revision 10)

---

# Overview

This document covers the following topics:

<i>Upgrade Notes</i> .....	3
<i>Caveats</i> .....	6
<i>Known Issues for 11.1.0</i> .....	7
<i>New Features in 11.1.0</i> .....	8
<i>Issues Addressed in 11.1.0</i> .....	19
<i>SL1 Extended Architecture</i> .....	30
<i>PowerPacks in 11.1.0</i> .....	31
<i>Disabling the Knowledge Base</i> .....	36
<i>Upgrade Process for Systems Running 8.1.0 and Earlier</i> .....	37
<i>Upgrade Process for Systems Running 8.1.1 and Later</i> .....	37
<i>Recently Deprecated Features</i> .....	38

# Upgrade Notes

**WARNING:** 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:

- 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

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.

**CAUTION:** There are issues with dynamic Device Groups and Business Services in new SL1 systems installed with the 11.1.0 ISO.

- If you are installing a new SL1 system from the 11.1.0 ISO or are using the 11.1.0 ISO to add new appliances to your SL1 system
- And you use or plan to use dynamic Device Groups or Business Services

You must immediately install the 11.1.1.2 upgrade after installing 11.1.0.

For details on the issue with dynamic Device Groups, see <https://support.sciencelogic.com/s/article/7546>.

**CAUTION:** If you are using one of the following PowerPacks, you must upgrade the to the specified version before upgrading to SL1 11.1.0. The previous versions will not prevent SL1 11.1.0 from installing or operating. But the previous versions might not collect data due to technical incompatibilities.

- Datacenter Advanced Enrichment Actions v106
- HTTP Action Type v103
- Linux SSH Automation v104

- Windows PowerShell Automation v104
- VMware Automation v102
- IBM: DB2 v103
- Linux: Base Pack v105
- NetApp: Base Pack v106
- Oracle: Database v103
- SL1: Concurrent PowerShell Monitor v102

**CAUTION:** There is a Known Issue with upgrades to 11.1.0. During the process to upgrade MariaDB with the `module_upgrade_mariadb` script, SL1 displays the error message "One or more config parameters missing in mysql.siteconfig". This is a spurious error that you can ignore. The workaround is to go to the **Possible next steps** option and select **Skip this error and proceed with upgrade re-run with -s option enabled**. For details, see the Knowledge Base article <https://support.sciencelogic.com/s/article/7091>. (JIRA ID: EM-47749) (JIRA ID: EM-47767) (Case ID: 00217851) (Case ID: 00219001)

- Because of major updates that were included in previous SL1 releases that you must consume before you can upgrade to 11.x, 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:

- 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
- During an upgrade to 11.1.0, the user interface will be unavailable due to an upgrade from PHP version 5 to PHP version 7.
  - **If you are running a version prior to 8.12.0, do not install 11.1.0 if you do not plan to immediately consume 11.1.0.** After you import the 11.1.0 release, all appliances in your SL1 system will now use the **new system update**. After you import the 11.1.0 release, you will not be able to stage and deploy any versions of SL1 previous to 8.12.0 or apply patches to versions of SL1 previous to 8.12.0. For details on the new system update, see the release notes for 8.12.0.
  - **To install 11.1.0 and its new System Updates tool, you must have already imported, staged, and deployed 8.6.0 or a later release.**
  - 11.1.0 includes important security updates. **You must reboot all AIO appliances and all appliances in the SL1 Distributed Architecture after deploying 11.1.0.** If you would like assistance in planning an upgrade

path that meets your security needs while minimizing downtime, please contact your Customer Success Manager.

- ScienceLogic strongly recommends that you review the [Known Issues](https://support.sciencelogic.com/s/topic/OTO0z000000E6w7GAC/known-issues) for SL1 (<https://support.sciencelogic.com/s/topic/OTO0z000000E6w7GAC/known-issues>) before installing a new update. For Known Issues specific to this release, see the [Known Issues](#) section of these release notes
- ScienceLogic strongly recommends that you review the instructions on planning an upgrade, best practices for upgrades, and executing an upgrade. To do so, see the chapter on **Upgrading SL1** in the [System Administration](#) manual.
- The following table specifies which SL1 updates require you to reboot all SL1 appliances and which SL1 updates require you to upgrade MariaDB.
  - Some SL1 updates include security updates. After applying these SL1 updates, you must reboot all SL1 appliances to apply the security updates. For instructions on rebooting, see the chapter on **Upgrading SL1** in the [System Administration](#) manual.
  - Some SL1 updates include an upgrade to MariaDB. These SL1 updates will automatically update MariaDB-client, MariaDB-common, and MariaDB-shared RPMs but will not update the MariaDB Server RPM. You must update the MariaDB Server RPM after you install the SL1 update. For instructions on updating MariaDB, see the chapter on **Upgrading SL1** in the [System Administration](#) manual.
  - SL1 updates are delta-less, meaning you install a single SL1 update file, and that SL1 update file can apply all SL1 updates between 8.6.0 and the current SL1 update, as needed. However, you might be required to reboot all SL1 appliances if one of the interim SL1 updates included a security update. And you might be required to upgrade MariaDB to the latest version if one of the interim SL1 updates included an upgrade to MariaDB.

For example, if you upgrade from SL1 8.12.2 to SL1 10.1.4, you will install only a single update. And looking at the table below, you can see that 10.1.4 requires you to reboot all SL1 appliances after upgrade but does not require you to upgrade MariaDB. However, two of the releases between 8.12.2 and 10.1.4 include an upgrade to MariaDB. Therefore, you must reboot all SL1 appliances and upgrade MariaDB after you upgrade to 10.1.4.

SL1 Release	Requires Reboot?	Required MariaDB Version
8.14.0	Yes	10.1.41
8.14.1	No	10.1.41
8.14.2	No	10.1.41
8.14.3	Yes	10.1.41
8.14.4	No	10.1.41
8.14.5	No	10.1.41
8.14.6	Yes	10.1.41
8.14.7	No	10.1.41
8.14.8	Yes	10.1.46
8.14.9	Yes	10.1.46
10.1.0	Yes	10.4.12
10.1.1	Yes	10.4.12
10.1.2	No	10.4.12
10.1.3	Yes	10.4.12
10.1.4	Yes	10.4.12
10.1.5	Yes	10.4.12
10.2.0	Yes	10.4.18
10.2.1	Yes	10.4.18
10.2.2	Yes	10.4.18
11.1.0	Yes	10.4.20

## Caveats

Consider the following caveats before deploying 11.1.0 :

- In SL1, all PHP code was converted to PHP7. If you have created custom content in PHP, see this page for backward compatibility: <https://www.php.net/manual/en/migration70.incompatible.php>
- Due to the PHP upgrade in 11.1.0, versions of Global Manager prior to 11.1.0 will not work with SL1 11.1.0 or Global Manager 11.1.0.
- Due to the PHP upgrade in 11.1.0, Web Proxy Services will not work in 11.1.0.
- PowerPacks built in SL1 version 11.1.0 and higher cannot be imported into previous versions of SL1 due to the upgrade to PHP version 7.0. However, PowerPacks built in previous releases of SL1 can be imported into version 11.1.0

- As of 10.1.0, SL1 no longer includes Flash.
- As of 8.12.2, ScienceLogic no longer updates the help that appears when you select the **[Guide]** button. The Unified UI provides a new tool for inline help. Under the user name in the upper right corner, click the down arrow and select **Help**. The browser will open a new window that displays the appropriate page from [docs.sciencelogic.com](https://docs.sciencelogic.com).
- As of January 1, 2021, **new installations of SL1 Extended Architecture are available only on SaaS deployments**. For existing on-premises deployments of SL1 Extended Architecture, please contact ScienceLogic Customer Support for upgrade documentation and help with technical issues.
- **SL1 updates overwrite changes to the configuration file `/opt/em7/nextui/nextui.env`**. This is a known issue (see <https://support.sciencelogic.com/s/article/11161> and <https://support.sciencelogic.com/s/article/1423>). ScienceLogic recommends that you backup this file before applying an update and then re-apply your changes to this file.
- 8.10.0 and later releases do 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 11.1.0.**
- The Knowledge Base includes known vulnerabilities for cross-site scripting and SQL injection. **If your first installation of SL1 was 8.9.1 or earlier, ScienceLogic strongly recommends that you [disable the Knowledge Base](#).**
- SL1 Extended Architecture does not provide MUD support.
- 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 ScienceLogic releases during a maintenance window.
- The ScienceLogic system update process starts a background process that can perform any required post-update 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-update 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.

---

## Known Issues for 11.1.0

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

- To configure CAC, user must define an LDAP service account with permissions that allow the service account to query LDAP. (Case: 00207174) (JIRA ID: EM-46927)
- The proxy configuration for Collector Pipeline does not support URLs that use HTTPS. This will be fixed in a future release.

---

# New Features in 11.1.0

## Admin Notifier

- The Admin Notifier tools now includes a new button, "View Events". This button displays the Events page, filtered to display only Admin Notifier events.

## Agent

- Updated SL1 logging to show aggregated Dynamic Application logging when a Dynamic Application is aligned with an agent.
- The default value for the Upload Interval on the Settings tab of the Device Investigator page was updated from 20 seconds to 1 minute. This update will result in substantial performance and stability improvements to the SL1 platform. This default value is enabled in version 174 and later of the Linux agent and version 133 of the Windows agent.
- A number of improvements were made to the Linux Agent installation process.
- The Linux agent installation program checks for the relevant version of libcurl before installing the agent. If libcurl is not installed, the agent installation program installs the most recent version of libcurl.
- The Linux agent was updated to work in a network with a proxy server by routing all Internet traffic through the proxy server. If a proxy server requires authentication, you will need to provide a username and password before traffic is allowed through the proxy server and to the Internet.
- The Windows and Linux agents were updated to work in a network with a proxy server, by routing all Internet traffic through the proxy server. If a proxy server requires authentication, you will need to provide a username and password before traffic is allowed through the proxy server and to the Internet.
- The RPM and DEB packages for Linux are now signed for added security.
- Agent users can now upload a data summary every sixty seconds using version 174 or later of the Linux agent.
- A number of improvements were made to the Linux Agent installation process.
- You can enable Run Book Actions and Run Book Automations to run with the agent by customizing the Run Book Automations and Actions in the following PowerPacks: - Linux SSH Automations - Windows PowerShell Automations version 103 or later.
- When creating or editing a Log File Monitoring Policy, a new "Other" option now appears in the "Source" drop-down on the Log Monitoring Policy modal. If you select "Other", a "Description" field then displays, allowing you to type the name of the event log source type.
- The Agents page has been updated to include commands for installing Linux agents on devices running Red Hat and CentOS - OS Libs (64 bit).

## Anomaly Detection

- When selecting available metrics for machine learning, the selection modal now includes the Dynamic Application name, presentation name, and index for each selectable metric.



- If you have enabled Machine Learning, SL1 default behavior is to publish CPU Vitals collected by Dynamic Applications

## Asset Management

- Added a new field for Model Number to the asset database.

## Business Services

- The sunburst chart is now compatible with N-tiers Services on the Overview page. Previously, this chart was only available for business, IT, and device services.
- For Business Services, add fields for Contact User and Contact Organizations, and RCA Options to the Service Investigator, in the Info drawer.
- In Business Services, added an information icon that appears if the number of devices/services is less than or equal to the maximum value allowed in a Business Service.
- A new Business Services Policy page was added.
- Users can now select more than one type of service in Services dashboard widgets, with the exception of Services overview and Sunburst visualization. It is no longer mandatory to select a Service type to create a Service widget.
- In Business Services, the Overview and Services/Devices tabs now indicate the current number of services/devices as well as the maximum number allowed.
- When clicking through a service map view or sunburst chart in Business Services, all of the major page components now reflect the details of the service selected from the hierarchy. These service details are reflected in the page title, Info drawer, and navigation tabs.
- Added the ability to bulk delete Business Service templates.
- In Business Services, the Map view has been updated to display only the currently selected service and its constituents.
- Search functionality in Business Services has been updated to enhance the usage of "and" and "or" within the advanced search syntax.

## CAC

- CAC authentication is now supported for logging in to the new user interface ("AP2"). CAC authentication in the new user interface ("AP2") has the same requirements as CAC authentication for the classic user interface. For more information about setting up CAC authentication, see the System Administration manual.

## Collection

- Users can now enable and disable concurrent PowerShell collection in the Collector Group Management and Behavior Settings pages. In 11.1.0, concurrent PowerShell collection is disabled by default.
- PowerShell concurrent Collection now includes CLI parameters to hide copyright banners, not load the PowerShell profile, or not present an interactive prompt to users.

- Users can now enable and disable concurrent SNMP collection in the Collector Group Management and Behavior Settings pages. In 11.1.0, concurrent SNMP collection is disabled by default.
- Users can now enable and disable concurrent SNMP collection for network interfaces in the Collector Group Management and Behavior Settings pages. In 11.1.0, concurrent SNMP collection for network interfaces is disabled by default.
- In 11.1.0, concurrent SSH collection is disabled by default. To enable it:
  - Go to the Process Manager page (System > Settings > Processes).
  - Find the process "Data Collection: SSH Collector" and select its wrench icon.
  - Set the Operating State to Enabled.
  - Save your change.

## Collector Pipeline

- SL1 administrators who use Collector Pipeline can now configure Streamer Push (on the Data Collector side) to communicate to Streamer (on the Compute Cluster side) through a proxy when direct line-of-sight is not available.
- Improved encryption between streamer push (runs on each Data Collector) and Database Server.
- Containerized services in Collector Pipeline now include hardened containers to improve security.

## Command Line

- "Previous methods for editing silo.conf, silo\_mysql.cnf, and firewall configurations have been deprecated.

The following commands are now the only supported methods for editing these configuration files:

File	Editing Tool
/etc/silo.conf	visilo
/etc/my.cnf.d/silo_mysql.cnf	vimariadb
/etc/siteconfig/firewalld-rich-rules.siteconfig	vifirewalld

The new commands ensure that changes are properly stored in their respective siteconfig file and will perform basic syntax checking of the file.

The following generator scripts have also been deprecated:

- /opt/em7/share/scripts/generate-silo-conf.py
- /opt/em7/share/scripts/generate-my-conf.py"

## Credentials

- You can now duplicate an existing credential (Save As) from the Credentials page.
- Updated the layouts of the Aliyun, AWS, Azure, Citrix Xen, and IBM universal credential types.

- On the Credentials page (Manage > Credentials), you can create a VMware credential type on the Create New drop-down. You can also create a VMware credential type during Guided Discovery for VMware.
- The Credentials page has been updated to automatically save any changes you make to the columns that appear on the page. When you move, sort, add, or remove columns from the page, those changes will be recalled the next time you visit the page.

## Dashboards

- New Organization widget for dashboards in the new user interface ("AP2").
- New Vitals Leaderboard in dashboards in the new user interface ("AP2").
- Interface Table widget, Leaderboard widget, and Leaderboard Bar widget in the new user interface ("AP2") support interface tags.
- In dashboards in the new user interface ("AP2"), all line charts for devices, interface, and file systems display the fetch count.
- In device, interface, and file system dashboards, you can now create a Leaderboard widget that contains table columns. This new feature enables you to see the top-n of certain metrics that you specify without calculating top-n values for the remaining metrics that are unnecessary, thus optimizing performance.
- To improve performance, newly created dashboards will have a default timespan of the Last 6 hours, rather than the previous default value of Last 24 hours. As always, you can manually change this timespan to fit your needs.
- You can now add an Organization Table widget to a dashboard. This widget displays information about 1 or more selected organizations, and can drive context to device, service, event, file system, and interface widgets.

## Device Groups

- Improvements to performance of dynamic Device Groups.

## Device Investigator

- Collection Labels were added to the drop-down list of metrics that you can add to the Device Investigator layout.
- The Ticket External Reference column values in the Events tab on the Device Detail page and Device Investigator now display as links to the external tickets.
- Pagination capabilities were added to the Cleared Events table that appears on the Events tab of the Device Investigator.

## Devices

- The Device Properties page now displays the Collector and Collector Group associated with a device.
- Custom links are now enabled by default in the new user interface ("AP2") for in the device drawer and in Device Investigator tools.

## Discovery

- A new IP Editor field type that end users can use to validate IP addresses was added to the Guided Discovery Framework.
- Guided Discovery Sessions are now saved to the database, and users can retrieve details of saved Guided Discovery Sessions with the `guidedDiscoverySessions` GraphQL query.
- A new Guided Discovery workflow, Ping, lets you create a new Pingable device at a given IP address.

## Events

- Custom links are now enabled by default in the new user interface ("AP2") in the events drawer.

## Global Manager

- Upon login, Global View is enabled by default and users see events from the child stacks.
- Users can access the Activity Center tools for devices and events from Global Manager.
- You can use GQL queries to retrieve the version and build number of the Global Manager system.
- Because they are not supported on Global Manager systems, the options to create Services and Maps dashboard widgets have been removed from Global Manager, while Organization widgets can be created only when Global View is disabled.
- In Global Manager systems, two new optional columns can now appear on the Devices and Events pages when Global View is toggled on. The columns, "Stack" and "Stack ID", respectively display the name and ID number of the stack from which the devices or events originate.

## GraphQL

- Added a feature that helps SL1 troubleshoot customer problems. GQL responses can include a log ID that allows SL1 staff to associate GQL messages and log entries.
- GQL API endpoints have been added to allow queries to determine which features are in-use, as well as the metrics related to platform performance.
- Added a new GQL mutation that creates or updates an organization.
- Guided Discovery Sessions are now saved to the database, and users can retrieve details of saved Guided Discovery Sessions with the `guidedDiscoverySessions` GraphQL query. Five GQL mutations were added to the CollectorGroup resource:
  - createCollectorGroup
  - updateCollectorGroup
  - deleteCollectorGroup
  - addCollectorsToGroup
  - removeCollectorsFromGroup.
- Removed items related to spritesheets and icons. Removed these queries: (sprite, sprites, optimalSprites), mutations (createSprite, deleteSprite, removeObsoleteSprites, removeObsoleteIcons), typeDefs (Sprite,

SpriteConnection, SpriteEdge, SpriteSearch, SpriteOptimizationModes, IconSpriteRef) and fields (Icon.isTombstoned, Icon.tombstone, IconSearch.isTombstoned and IconSearch.tombstone). We no longer dynamically produce spritesheets from user provided icons. Icons queries and mutations are still available through GQL. Dynamic spritesheet support has been removed. Icons are now deleted at time of delete rather than being marked for deletion at a later time.

## Maps

- When designing maps, you can now specify whether the map nodes are represented by icons, images, or solid colors that correspond with the nodes' current status (healthy, notice, minor, major, or critical).
- Added the ability to reverse direction for topology relationships in Layer 2, Layer 3, CDP, LLDP, and ad hoc maps. To do so, click the relationship link, click the Change Direction icon on the Relationship pane, and then click the Save icon.
- Improved the user experience for maps that contain a large number of nodes, and increased the maximum number of nodes that can display on maps.

## Multi-Tenancy

- Collector groups can now be assigned to specific organizations. By default, collector groups are assigned to all organizations; however, administrators can override that setting and assign collector groups to one or more specific organizations. Non-administrator users can view and align only those credentials that are assigned to organizations common to both the user and the device's collector group, plus those credentials that are assigned to all organizations or otherwise required for that collector group. For more information, see the System Administration manual.
- You can align a collector group to one or more organizations.

## PhoneHome Collectors

- PhoneHome shell logs for PhoneHome users are stored in the user's home directory in the log subdirectory, which prevents issues that occur if multiple PhoneHome users are trying to write to the same log file at the same time. The shell log files are auto-rotated to help prevent logs from filling the home partition.
- Links were added to the PhoneHome Collectors wizard to download (or access in a cloud marketplace) an appropriate image of the SL1 collector, or access instructions for installing the collector.

## PHP

- In SL1, converted all PHP code to PHP7. If you have created custom content in PHP, see this page for backward compatibility: <https://www.php.net/manual/en/migration70.incompatible.php>
- Due to the PHP upgrade, versions of Global Manager prior to 10.2.0 will not work in 11.1.0.
- Due to the PHP upgrade, Web Proxy Services will not work in 11.1.0.
- PowerPacks built in SL1 version 11.1.0 and higher cannot be imported into previous versions of SL1 due to the upgrade to PHP version 7.0. However, PowerPacks built in previous releases of SL1 can be imported into version 11.1.0
- During an upgrade to 11.1.0, the user interface will be unavailable due to an upgrade from PHP version 5 to PHP version 7.

- In SL1, converted all PHP code to PHP version 7. For more information about PHP 7.0.0, see [https://www.php.net/releases/7\\_0\\_0.php](https://www.php.net/releases/7_0_0.php).
- SL1 supports upgrades from SL1 versions that use PHP version 5 to a version that uses PHP version 7.
- The OneLogin SAML library has been upgraded to v3.5.1 for compatibility with PHP versions greater than 7.1.
- SL1 version 11.1.0 includes PHP version 7.4.21. When you upgrade to SL1 version 11.1.0, the PHP version is automatically upgraded for you. Older PHP versions are not supported. For more information about PHP 7.x, see [https://www.php.net/releases/7\\_0\\_0.php](https://www.php.net/releases/7_0_0.php). Because of these changes, classic Global Manager is no longer supported in SL1 version 11.1.0 or later.

## Platform

- Address selinux for docker containers
- The system status script has been enhanced for this release to collect a wider variety and broader scope of logs and to check for broken symlinks in /data.local.
- The system status script will now include the upgrade\_mariadb.log.
- SL1 will now throttle surges of multi-index Dynamic Application alerts on individual devices and will notify the system administrator when alerts for a device exceed the threshold.
- The default innodb buffer pool size on Database Servers is now adjusted based on memory size.
- The timeout value has been increased for MariaDB start and stop.
- Upgraded Node.js to node 14.
- You can now put the Database Server in maintenance mode and stop all pull processes from the Data Collectors. You can then perform database maintenance or network maintenance without generating events. The new commands are silostart and silostop. )

## PowerFlow

- A new Change Events tab is now available for the Service Overview page for PowerFlow users. This tab displays information about the events that are created when PowerFlow pulls change data from ServiceNow, including both active and cleared events. (NOTE: To enable and use this tab, users should contact ScienceLogic Professional Services.)

## Reports

- Updated EM7 Core Reports PowerPack (Version 115 and Revision 5056) included with SL1 11.1.0.
- You can generate reports during the licensing grace period (90 days after license expires).
- Leaderboard metrics have a new field called "Leaderboard from:" with the following options available in a drop-down list: "Full timespan" and "Last poll".
- Tooltips for Gauge Chart, Pie Chart, Bar Chart, and Number visualizations, Tables, and Leaderboards have been updated to show the time of data collection for Device, Interface, Filesystem, and DCM widgets, if available.

## Run Book Automation

- When an SL1 environment has less than 20 Run Book Automation (RBA) executors and its Database Server's maximum database connections are equal to or greater than 500, the total number of RBA executors will be set to 20.
- You can now bulk enable or disable multiple automation policies from the Run Book Automation Policy Manager page.

## Scheduler

- Updated the SL1 Job Scheduler to properly handle transitions into and out of Daylight Savings Time so that the execution of any schedule does not take place one hour earlier or one hour later than expected.

## ScienceLogic API

- Added the following new endpoints to the ScienceLogic API to support configuration of the Streamer Push proxy:
  - GET /sladmin/v1.0/streamerpush/proxy (Returns the current proxy configuration information)
  - POST /sladmin/v1.0/streamerpush/proxy (Allows you to set the proxy information)
  - POST /sladmin/v1.0/streamerpush/proxy/toggle (Allows you to toggle proxy on or off without deleting the configuration)
- The following actions were added to the REST API endpoint "collector\_group": - all\_orgs - aligned\_organizations To enable multi-tenancy for collector groups, the database setting "master.system\_settings\_core.enable\_cug\_orgs" must be set to 1. When multi-tenancy is enabled, an administrative user can update all collector groups using the new fields. Non-administrative users can update all collector groups for which the "all\_orgs" field is set to 1. Otherwise, these users can only update credentials and collector groups within their aligned organizations.
- For API fields that require an account URI, such as /api/account/2 as the supplied value in filters when searching or in POST bodies for write operations, you can instead use /api/account/\_self as a shorthand for the URI of the currently logged-in user.

## Security

- Upon installation of 11.1.0 or upgrade to 11.1.0, SL1 defaults to HTTPS. You can change this setting in the Behavior Settings page (System > Settings > Behavior).
- SL1 administrators can display a block of text that users must accept every time they log in to SL1.
- The containerized service that allows SL1 to use multi-processing to collect data from SNMP Dynamic Applications (released with 10.1.0) now includes hardened containers to improve security.
- The containerized services in Collector Pipeline now include hardened containers to improve security.
- Stronger encryption for credentials.

- On the Discovery Session Editor page of the classic user interface, in the SNMP Credentials section, the Other Credentials section, and the Collection Server drop-down are now aware of the organization and will display only credentials or collectors available to the selected organization.
- The SNMP read and write credentials and the collector group drop-downs are limited only by the settings of "all\_orgs" and "aligned\_organizations". The secondary credential shows only credentials aligned to the device.
- You can select the "Include PowerPack Sensitive Fields" on the Behavior Setting page (System > Settings > Behavior) page if you want to include sensitive fields when sharing the PowerPack. These sensitive fields include passwords and SSH keys.
- Two new role-based accounts, sl1 user and sl1 admin, are included in this release. These accounts allow you to provide limited system access to certain commands on the SL1 appliance or troubleshoot system issues. To use either of these accounts, you must log into the appliance through Secure Shell (SSH). All actions are logged, and an SL1 Administration can monitor and take control of the session if necessary. These accounts do not appear in the User Accounts page and cannot be administered through the UI. For details, see the manual *Organizations and Users*.
  - The "sl1 user" role-based account has a limited menu of options available through the command-line interface (CLI) and can be used by remote operations personnel to perform basic tasks.
  - If you are using the sl1 admin role-based account, you can obtain the one-time password on the Appliance Manager page (System > Settings > Appliances). To obtain the password, click the "Get One Time Password" icon next to the appliance you are accessing, and enter the 3- or 9-digit code you received at the command line.
  - You can now grant a one-time use password to a role-based administrator account, "sl1 admin". This account allows you to permit less-trusted personnel (for example, contract personnel) to perform a limited set of administrative commands during a single session.
- If a user tries to access a page via URI and the user does not have permissions for that page, SL1 denies access and creates an audit log entry. That page does not show up in the user interface if the user does not have the right permission for that page.
- When users create, edit, or delete a Dynamic Application, SL1 creates an audit log message.
- The containerized services that run on Data Collectors now includes hardened containers to improve security.
- The containerized service that processes data from Dynamic Applications (released with 8.14.0) now includes hardened containers to improve security.
- Streamer push (the containerized service that runs on the Data Collector and sends collected data to the Database Server or the API endpoint) now includes hardened containers to improve security.
- The containerized service that allows SL1 to use multi-processing to collect data from SNMP Dynamic Applications (released with 10.1.0) now includes hardened containers to improve security.
- To enable a Content Security Policy (CSP), add the following environment variables to the /opt/em7/nextui/nextui.envfile file: INCLUDE\_CONTENT\_SECURITY\_POLICY\_HEADER=enabled  
CONTENT\_SECURITY\_POLICY\_HEADER=default-src 'self';



## SNMP Traps

- A new SNMPv3 Trap Configuration Reset button now appears at the top of the Credentials page (Manage > Credentials). When you click this button, SL1 automatically configures the `/etc/snmp/snmptrapd.conf` file on Data Collectors and Message Collectors, so Data Collectors and Message Collectors can accept traps from monitored devices and communicate with those monitored devices. To use this feature, users must have the `SYS_SNMP_TRAP_CONFIG_PUSH` Access Hook aligned to their user profiles.

## Subscription Billing

- The classic dashboards, Dynamic Application, and widget definitions have been removed from the ScienceLogic Subscription Audit Pack, which is included by default in the 11.1 platform release. The Manage > Subscription Usage page displays all the data that was on the classic dashboard. This content is not automatically deleted when a 10.2 or earlier system is upgraded to 11.1. This content will continue to function but will no longer receive maintenance updates.
- A new Dynamic Application, Support: Feature Usage Configuration, has been added to the ScienceLogic Support PowerPack. This Dynamic Application indicates which primary features of SL1 are in-use. .
- The device template for SL1 database servers has been updated to align this new Dynamic Application . The data collected by this Dynamic Application will be sent to ScienceLogic via the subscription usage mechanism in 11.10 GA.
- Additional data retention thresholds for stored usage and telemetry data have been added.

## SSO

- A new "Request Signature Algorithm" field is available in the SSO Resource Editor when SSO authentication "Signing Options" field is set to "SP Request and IdP Response". Available options are: RSA-SHA256, RSA-SHA384, and RSA-SHA512.

## Subscription Billing

- The Subscription License Total Storage Report has been updated to use the Available Space metric for measuring storage capacity on Solidfire devices instead of the Max Provisioned Space metric.
- The usage payloads sent to ScienceLogic now include CPU, Memory, and Disk configuration and usage information for SL1 appliances. The usage payloads sent to ScienceLogic now include the number of devices and interfaces monitored by each Data Collector.
- The following configuration data and metrics are included in the new "Feature Usage" tab for devices assigned an SL1 appliance device class: Devices per Collector, Interfaces per Collector, Total Storage Size on SL1 appliances, Used Percent from the Support: InnoDB Size Dynamic Application, Used Space from the Support: InnoDB Size Dynamic Application, Load Average from the Support: SL1 Performance Dynamic Application, CPU Overall Usage % from the Linux CPU Stats Dynamic Application (Linux Base Pack), Physical Memory Utilization from the Linux: Memory Stats Dynamic Application (Linux Base Pack), Number of Cores from the Linux: CPU Config Dynamic Application (Linux Base Pack), Total Physical Memory from the Linux: System Config Dynamic Application (Linux Base Pack) ^ Combo of PTEL-1040: Contract Compliance Data Reporting, PTEL-1041: Project Titan Feature Usage Data

- The payload of license and telemetry data sent to ScienceLogic now includes information about which SL1 features are in-use.
- The process that sends license and telemetry data to ScienceLogic has been updated for scale and performance. This includes splitting processing across multiple processes that run at different times during the day.

## Support PowerPack

- The "Support: File System Dynamic Appliance" discovery object in the Support PowerPack, included in SL1, changed to align only with SL1 appliances. Customers should un-align this discovery object from all non-SL1 appliances

## System Status Script

- The system status script will now check the number of data pull processes configured and, if less than the recommended default number, will print out an error while running.
- The system status script now displays a message if jemalloc is not the default malloc.

## System Update

- During a system update, a banner alert will show the number of appliances that will not be patched due to one of the following statuses: patch disabled, patch ineligible, or proc\_mgr paused.
- When an SL1 instance that was updated via AMI is rebooted, the hostname will be preserved.
- A new sub-command, "sysuptb check-extra-pkgs", was added to the sysuptb troubleshooting script. This command checks and lists any extraneous (non-SL1) packages installed on the appliance, which can help to identify a root cause for failures related to package dependency issues.
- The banner on the Updates page (System > Tools > Updates) will display the number of appliances that need a MariaDB server update.

## Themes

- In the EM7 Base Themes PowerPack and the SL1 Unified Theme PowerPack, improved the highlighting in the code view, for easier viewing.

## User Interface

- Upon installation of 11.1.0 or upgrade to 11.1.0, the new UI ("AP2") is the default UI. To set the UI to the classic UI, go to the Behavior Settings page (System > Settings > Behavior) and unselect "New UI as default". To change the UI to the classic UI while working, append "/em7" to the URL. To change the UI to the new UI ("AP2") while working, append "/" to the URL.
- Upon installation of 11.1.0 or upgrade to 11.1.0, the new UI ("AP2") is the default UI.
- Added the checkbox New UI Default in the Behavior Settings page (System > Settings > Behavior) that allows you to toggle the new UI ("AP2") as the default UI.

- When you hover over the user name in the upper right corner of SL1, the hover text will display the user's time zone.
- Removed flashed-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 SL1 page title appears in the title bar of your browser, and the page title also appears in your browsing history to help you navigate through the pages you have visited.
- The SL1 login button displays a spinning status icon while you are being logged in.
- SL1 now retains the search criteria on pages in the new user interface ("AP2") after you refresh the browser or navigate to another page and then return to the search page using the Back button in the browser. You can also bookmark this URL to start at this search or to share this search with another user.
- Updated the Single Instance Login settings so they can be configured to work in the new user interface ("AP2") in addition to the classic user interface. These settings specify whether more than one instance of a single username can be logged in to the user interface at the same time and are found on the Behavior Settings page (System > Settings > Behavior). To use the feature in the new user interface ("AP2"), you must log in to the SL1 appliance and edit the environment file `/opt/em7/nextui/nextui.env`; add `"#"` to the beginning of the line that includes `"AUTH_CACHE=300000"` to comment out the `AUTH_CACHE` variable; and then restart the server using the command `sudo systemctl restart nextui`. For more information, see the System Administration manual.

---

## Issues Addressed in 11.1.0

### Agent

- Addressed an issue in which agent\_discovery storage objects for a modeled P0 Agent device were failing when Auto-Update was disabled on the device. (Case: 00102448) (JIRA ID: EM-38633)
- Addressed an issue with agent log data where log lines were not in UTF-8 encoding.
- Addressed an issue where the SL1 Linux agent v178 generated false-positive alerts. (Case: 00191605) (JIRA ID: AP-1950)
- Addressed an issue where the Linux agent would time out while trying to upload data, and the data did not get uploaded to SL1. (Case: 00179027) (JIRA ID: EM-45356)
- Addressed an issue where regular expressions that exceeded 64 characters were cut off and sent to the agent in an incomplete state. The character limits was increased to 255 characters. (Case: 00178211) (JIRA ID: EM-44652)

## Assets

- The Collections Objects page for Dynamic Applications of type "Configuration" now includes a **Precedence** drop-down list for **Asset Link**. This field allows you to define the precedence when multiple data sources are defined for an asset link. All asset fields have the default priority value set to 50, but you can adjust the value manually. The asset field with the highest precedence that also contains valid data is visible in Asset Properties. (Case: 00098398, 00136046) (JIRA ID: EM-29177)

## Business Services

- Addressed an issue that prevented the Metrics drop-down values from loading when editing status policies in services. (Case: 00144693) (JIRA ID: EM-42175)
- Addressed an issue that was preventing the advanced search feature in business services from returning results when filtering on health, availability, and risk in environments using MariaDB. (Case: 00154338) (JIRA ID: EM-41992) (JIRA ID: EM-42134)
- Addressed an issue in Business Services that was preventing users from successfully searching for devices when the search query included both custom attributes and event messages. With this change, you can now include both custom attributes and event messages in a successful device search. (Case: 00130912) (JIRA ID: EM-40202)(EM-41527)

## Classic Maps (Views )

- Added a new feature to the Views pages. For customers using HTML5 maps in the classic UI, users can specify whether relationships between devices should use the status color of the device on each end or the status color of the network interface on each end. (JIRA ID: SLUI-10169)
- Added a new feature to the Views pages. For customers using HTML5 maps in the classic UI, users can specify whether relationships between devices should use the status color of the device on each end or the status color of the network interface on each end. (JIRA ID: SLUI-9800)

## Concurrent SNMP Collection

- Concurrent SNMP collection no longer drops polls and creates data gaps when querying SNMPv3 devices with large latency. (Case: 00187707) (Case: 00205232) (JIRA ID: EM-45866) (JIRA ID: EM-47111) (JIRA ID: EM-45093)

## Content Library

- Addressed an issue in which the association of a content library with an execution environment did not synchronize the same way when performed from the UI versus the REST API. Both the UI and REST API now synchronize the underlying database model the same way. (JIRA ID: EM-39947) (JIRA ID: EM-39784)

## Credentials

- Encryption standards for new and existing SSH credentials have been strengthened. (Case: 00157192) (JIRA ID: EM-15226) (JIRA ID: EM-15227) (JIRA ID: EM-42569) (JIRA ID: EM-43663)

## Custom Attributes

- Addressed an issue that was preventing users from properly saving multiple extended custom attributes when one was a string type and another was an integer type. (Case: 00144228) (JIRA ID: EM-41387) (JIRA ID: EM-41303)

## Dashboards

- Addressed an issue in dashboards that was preventing Map widgets from displaying properly when the dashboard also contained a Leaderboard widget. With this update, dashboards no longer exhibit sporadic issues displaying Map widgets and Leaderboard widgets at the same time. (Case: 00148795) (JIRA ID: EM-41772) (JIRA ID: SLUI-9232)
- Addressed an issue in which users could not create interface line chart widgets without driving them from a leaderboard or table widget in dashboards. With this update, interface line chart widgets no longer require the "Interfaces can be selected from other widgets" option to be enabled, thus allowing users to create standalone interface line charts in dashboards. (Case: 00146282) (JIRA ID: EM-41510)
- Addressed an issue that was preventing some shared and public dashboards that should have been visible to a user from displaying in the list of dashboards when there were a large number of dashboards that the user did not have permission to view at the top of the dashboard sort order. (Case: 00135171) (JIRA ID: EM-41184)
- Addressed an issue in the classic user interface that was causing fewer than 10 devices to display in Top 10 dashboard widgets. (Case: 00133335) (JIRA ID: EM-41136)
- Addressed an issue with filtering Business Services in dashboard widgets that was causing the widgets to not load properly. (Case: 00161978) (JIRA ID: EM-43184) (JIRA ID: EM-43343)

## Devices

- The reports in the Device Performance page (Registry > Devices > Device Manager > bar-graph icon > Performance tab) no longer display the incorrect time or a time greater than 24:00:00. (Case: 00161436) (JIRA ID: EM-43592) (JIRA ID: EM-43419)
- Addressed an issue where devices with SSL certificates expiring after January 19th, 2038, are inaccurately reported as "expired". (Case: 00053541, 00081430) (JIRA ID: EM-42877, EM-32807)
- Addressed an issue that was causing the Device Investigator Overview panel to improperly scale when encountering long IPv6 addresses. With this update, the panel now better handles long titles and IPv6 addresses. (Case: 00155346) (JIRA ID: EM-42381)
- Addressed an issue that was preventing a device's cleared events from displaying in the Device Summary in the new user interface ("AP2"), due to a large number of cleared events. With this update, the Device Events page now displays 30 cleared events at a time to avoid this issue when a device has a large number of cleared events. (Case: 00137118) (JIRA ID: EM-40698)
- Addressed an issue where the event policy for "maintenance window opening" used an expiration value that overrode the actual status of a device, specifically when that device remained in maintenance because of overlapping maintenance schedules for the device. (Case: 00109304) (JIRA ID: EM-38746)
- Addressed an issue where results from the SNMPDump tool within the Device toolbox were showing stderr logs in addition to standard logging.

- Device Management Addressed an issue where polling continues on a device during a maintenance schedule. When collection is set to be disabled in a maintenance schedule for a device, it will be disabled for the entire period of the maintenance schedule, unless a patch window was also set in the schedule. When a patch window is configured for a maintenance schedule, collection is only disabled during the patch window, which is typically triggered by a reboot of the device. The "Collection Polling" section of the Device Management manual was updated with this information. (Case 00004073) (Support ID: 170560) (JIRA ID EM-26452)
- The System Vitals Summary report for a device (Devices > Device Manager > bar-graph icon > Performance tab) now displays the correct legend. (JIRA ID: EM-041941) (Support ID 27388) (Case: 00098186)
- The ARP Ping tool in the Device Tools page (Devices > Device Investigatory > Tools) or (Registry > Devices > Device Manager > wrench icon > Toolbox) no longer generates sudo permission errors. (Case: 00039376) (JIRA ID: EM-40673)
- Addressed an issue in which non-ASCII characters in a device name or description could trigger unhandled exceptions. (Case: 00185241) (JIRA ID: EM-45068)

## Device Groups

- Reduced the amount of time it takes SL1 to evaluate Dynamic Device Group rules, which in previous versions caused SL1 to refresh the membership of the device group too slowly, resulting in events not being suppressed for devices in the device group. (Case: 00107527) (JIRA ID: EM-39070)
- Addressed an issue in which event suppression was failing for dynamic device groups. (Case: 00057562) (JIRA ID: EM-34516)

## Device Templates

- Addressed an issue where the Service Policy tab of the Device Template Editor modal was empty when there were 500 or more policies to load into the modal for selection. (Case: 00062089, 00084687, 00099673, 00157890, 00172522) (JIRA ID: EM-35276)
- Addressed an issue where collection objects that were disabled through a device template would get enabled again after nightly discovery. (Case: 00093155) (JIRA ID: EM-37419)

## Discovery

- Addressed an issue that was causing the pointer (PTR) record to be stored instead of the fully qualified domain name (FQDN) when the Enable DHCP option was selected during a discovery session. (Case 00047649) (JIRA ID: EM-32143)
- Addressed an issue that was causing availability checks to be disabled on devices that were discovered with DHCP enabled. Internal collections will now be enabled when devices are discovered via hostname/DNS. (Case: 00003556) (Case: 00036874) (Support ID: 133293) (JIRA ID: EM-20352)

## Documentation

- Fixed an error in the **Subscription Billing** manual. Added the full path for silo.conf. (Case: 00195259) (JIRA ID: EM-45643)

## Dynamic Applications

- When a PowerShell Configuration Dynamic Application or PowerShell Performance Dynamic Application is aligned to a device with a non-PowerShell credential, SL1 no longer generates an unhandled exception occurred and no longer stops collection for all Dynamic Applications that were to be collected in that polling interval. (JIRA ID: EM-42219) (Case: 00166623)
- Descriptions for SNMP Dynamic Applications no longer display as HEX. Fixed a processing problem with null-terminated strings, (JIRA ID: EM-42689) (JIRA ID: EM-41417) (Case: 00143289)
- When you execute a cache-producing Dynamic Application, that Dynamic Application produces a single storage object rather than two. (JIRA ID: EM-41053) (JIRA ID: EM-41049) (Case: 00137968)
- Data with negative values from Performance Dynamic Applications is now normalized correctly. (JIRA ID: EM-40930) (JIRA ID: EM-40600) (Case: 00054638) (Support ID: 116615) (JIRA ID: EM-17070)
- Addressed an issue where aligning a performance Dynamic Application without a presentation object failed, but the user interface indicated that the performance Dynamic Application was successfully aligned. (Case: 00054638) (Support ID: 116615) (JIRA ID: EM-17070)
- When a collection object in a Dynamic Application fails to collect a value for a specified amount of time, the resulting empty string will no longer disable the collection object. (Case: 00106497) (JIRA ID: EM-43633)

## Email

- In the Email Settings page (System > Settings > Email), the **Authorized Email Domains** field will no longer accept domains longer than 128 characters. (Case: 00099022) (JIRA ID: EM-43255)

## Events

- Updated the default sorting options in the Event Console to sort events by Severity and then Last Detected, to ensure that the highest severity and most recent events appear at the top of the list. If users change their sort options, SL1 will remember those changes and sort events in that same manner the next time they visit the page. (Case: 00153240) (JIRA ID: EM-42159)
- Addressed an issue related to Entity Name displayed in the Event Console for an event record; Entity name will now match text capitalization for that entity as it is registered in its respective table.
- Addressed an issue in which "certification expired" events were being incorrectly auto-cleared by events relating to other certifications. With this update, certification-related events will auto-clear only when the certification ID is the same. (Case: 00026110) (JIRA ID: EM-30159)
- Addressed an issue where two users click "Create Ticket" or the "life ring" icon for an event at the same time, and one of the two tickets being updated is not saved as a result. (Case: 00098190) (Support ID: 99713) (JIRA ID: EM-14098)
- In the Event Policy Editor page (Registry > Events > Event Manager > create button or wrench icon ) or (Advanced Menu > Events > Event Policies > Create/Edit > Match Logic), for events with an **Event Source** of *Trap*, the **Source Host Varbind** field ignores input case (uppercase, lowercase, or a mixture of both). (Case: 00138920) (JIRA ID: EM-41959)

## Execution Environments

- When multiple versions of the same library are added to an execution environment, the execution environment will load only the library with the highest version. (Case: 00140987) (JIRA ID: EM-43352)

## Global Manager

- Addressed an issue that was preventing users from acknowledging events in Global Manager if there was no user with the same user ID on the destination stack. With this update, SL1 no longer requires identical user IDs on the stacks in order to display the correct user that acknowledged an event in Global Manager. Additionally, a new acknowledgeEvents GraphQL mutation was introduced to the SL1 GQL schema to enable bulk event acknowledgment, including in Global Manager mode. (Case: 00156203) (JIRA ID: EM-43080)

## GraphQL

- Addressed an issue that was causing the "Monitored" value to always display as "True" when querying device processes in GraphQL, regardless of the actual value. With this update, the GQL query now returns the correct "Monitored" value. (Case: 00135641) (JIRA ID: EM-40738)
- Addressed an issue with the GraphQL schema that was causing an error when users attempted to return the percentUsed field when performing a fileSystemVitalIndexes query. The field has been updated as a "Float" type to prevent this error. (Case: 00116738) (JIRA ID: EM-39132)
- Addressed an issue that was causing the software GraphQL query to always return null values. With this update, the software query was removed from the schema. (Case: 00110313) (JIRA ID: EM-38718)
- GQL can now query the Type option for the Asset resource. (JIRA ID: EM-43534) (JIRA ID: 00157834) (Case: 00157834)

## High Availability and Disaster Recovery

- Addressed an issue where systemd tried to restart the NextUI service in a Disaster Recovery node. (Case: 00034551) (JIRA ID: EM-31700)
- The DRBD Proxy license check in coro\_install has been improved. The coro\_install will no longer fail if the DRDB Proxy license check fails. (Case: 00133885) (JIRA ID: EM-40437)
- CRM templates have been updated to prevent replication disconnection on DR failover. New templates must be manually applied. The system will alert you if your appliance can benefit from the changes. (Case: 00103667) (JIRA ID: EM-37806)
- Removed the option for node maintenance in coro\_config. Added cluster-wide maintenance option in coro\_config. (Case: 00104428) (JIRA ID: EM-37910)

## Installation

- The SL1 AMI now installs successfully on an AWS M5 Instance. (JIRA ID: DO-3857) (JIRA ID: EM-43034) (Case: 00160579)



## Internal Collections Dynamic Applications

- Internal Collections Dynamic Applications (ICDA) no longer cause false alerts for process inventory, services inventory, and file system inventory. (JIRA ID: EM-43783) (JIRA ID: EM-43553) (Case: 00161944) (Case: 00157364) (Case: 00164303) (Case: 00167935) (Case: 00175579) (Case: 00179779)

## Licensing

- The `licensed_state` command now displays an error when run as a non-root user, rather than providing a false state of being unlicensed.

## Logs

- Addressed an issue that was causing user credential passwords to display in plain text in developer logs. With this update, plain-text credentials no longer display in the logs. (Case: 00098701) (Support ID: 122701) (JIRA ID: EM-18787)
- Addressed an issue in which results from the SNMP Dump tool within the Device toolbox were showing stderr logs in addition to standard logging. (Case ID: 00095578) (JIRA ID: EM-37440)
- Addressed an issue in which the system status script would report certain log files as being owned by the wrong user or group. (Case: 00113405, 00136043, 00089958, 00137951) (JIRA ID: EM-38879, EM-40605, EM-36419)

## Machine Learning

- If you have enabled Machine Learning, SL1 now publishes CPU Vitals collected by Dynamic Applications by default. (JIRA ID: EM-43232) (JIRA ID: EM-43069)

## Maps

- Addressed an issue in which CDP maps were not displaying the correct parent-child relationships for some daisy-chained devices. With this update, users can manually adjust the direction of the parent-child relationship between devices. (Case: 00154621) (JIRA ID: SLUI-10862) (JIRA ID: EM-42271)

## Network Interfaces

- Addressed an issue that was causing unhandled exceptions to appear in the `silos.log` when `high_precision_if_collect` value is 1 and no data was seen in Performance Graphs for Device Network Interfaces Interface graphs. (Case: 00136061) (JIRA ID: EM-41045)
- Addressed an issue that was causing unhandled exceptions to appear when the `collect_if.py` had a "0" value. (Case: 00047094) (JIRA ID: EM-39687)
- Addressed an issue that was causing unhandled exceptions during nightly discovery when internal collections Dynamic Application for Interface Inventory is aligned to a device and the Auto-Update setting is disabled on the device. (Case: 00135932) (JIRA ID: EM-31488)
- If a network interface name includes a Latin-1 character, SL1 no longer stops collecting data for that interface and no longer stops the Interface Alert service. (Case: 00208855) (JIRA ID: EM-47112) (CPL-575)

## Organizations

- Addressed an issue where you could not edit a note using the HTML Editor on the Notes tab for an Organization. (Case: 00099209) (Support ID: 117020) (JIRA ID: EM-17237)

## Platform

- Addressed an issue that sometimes displayed an incorrect default status policy for a device service. (Case: 00156713) (JIRA ID: EM-43134)
- Audit rules now load as expected. Audit rules that had syntax errors, that were duplicates, or that targeted files that did not exist have been fixed.
- Addressed an issue where the licensing graphic on the System Usage page (System > Usage > System Usage) did not display, as the graphic used Flash. (Case: 00128285) (JIRA ID: EM-40177)
- Users who installed SL1 prior to the 8.8.0 will no longer experience significant increases in CPU utilization during nightly update Discovery. (JIRA ID: EM-38625) (JIRA ID: 33409)
- To improve performance, removed an index and added a primary key to the device\_services table. (JIRA ID: EM-40590) (Case: 00136001)
- Improved log rotation for /var/log/em7/snmp\_agent.log. (Case: 00072979) (JIRA ID: EM-35685)
- Made multiple improvements to the alert function "deviation ()". SL1 now supports deviation policies that move between Data Collectors during load balancing, The weeks\_collected metric that lives on Data Collectors is now accurate. (Case: 00164968) (JIRA ID: EM-44141) (JIRA ID: EM-43773)
- SL1 Administrators are now notified if the Enterprise Database: Collector Config Push process (config\_push.py) fails due to a corrupt database table on the Data Collector. (Case: 00144225) (JIRA ID: EM-41283)(JIRA ID: EM-39911)
- Moved the timeout setting for operating-system processes to the master.system\_custom\_config table, in the os\_process\_timeout column, instead of using the value in the credential. The timeout value has been increased to prevent multiple timeouts during upgrade. (Case: 00186868) (Case: 00181625) (JIRA ID: EM-46551) (JIRA ID: EM-46095) (JIRA ID: EM-47003) (JIRA ID: EM-45197) (JIRA ID: EM-45291)

## PowerPacks

- The "Support: DB Space Estimator" Dynamic Application in the Support PowerPack now returns more accurate estimates. (Case: 00137925) (JIRA ID: EM-41215)
- Code in the Cisco: ACI PowerPack has been refactored to improve performance and increase reliability during APIC failover by maintaining sessions with all APICs. (Case: 00096286 , 00116626; JIRA ID: SOL-10925)
- In the Cisco: ACI PowerPack, an issue was addressed in which the "Cisco: ACI Faults" was exhibiting unexpected behavior with fault filtering. New alerts have been implemented that are triggered when principal ACI nodes are discovered. (Case: 00097058, 00079453; JIRA ID: SOL-6899)
- The Cisco: API PowerPack was updated to allow direct upgrade from any version newer than 103 directly to 110. Previously, if you upgraded from any version older than 108, you may have needed to delete all devices and re-discover. (Case: 00065408; JIRA ID: SOL-2559)

- In the Cisco: Cloud Services Platform PowerPack, an issue was addressed in the "Cisco: CSP 2100 Services Discovery" and "Cisco: CSP 2100 Service Resource Stats" Dynamic Applications in which an exception was occurring when services were not configured to run on a device. (Case: 00059262; JIRA ID: SOL-1487)

## PowerShell

- Added additional details to the powershell\_collector.log to enable troubleshooting. The (JIRA ID: EM-43200) (Case: 00162182)

## Publisher

- SL1 Publisher now supports a data model for Dynamic Application performance data from SL1.

## Reports

- Addressed an issue where there was an unnecessary mailto-link in the reports sent via email when the username had an at (@) and dot (.). (Case: 000154580) (JIRA ID: EM-42065)
- Addressed an issue where a user with "emissary" rights to an interface report could not view the graphs in those reports, as those graphs still defaulted to using Flash. (Case: 00148891) (JIRA ID: EM-41573)
- Addressed an issue where the "Device Outage History" report did not accurately show how long a device was not available. (Case: 00124312) (JIRA ID: EM-40689)
- Addressed an issue where reports did not display the logo associated with the theme for that user. (Case: 00098622, 00156241) (Support ID: 143355) (JIRA ID: EM-21575)
- Addressed an issue where the "Report Span Workday With Timezones Always Avail" component selected the last day of the previous month by default instead of the first day of the current month. (Support ID: 122422) (JIRA ID: EM-18395)
- Scheduled reports (Reports > Create Report > Scheduler) with a defined output of HTML now include the entire report when emailed to recipients. (JIRA ID: EM-43963) (JIRA ID: EM-43679) (Case: 00165435)
- The Interface Usage Report (Reports > Run Reports > Network Interfaces > Interface Usage) now displays accurate data for % utilization. (JIRA ID: EM-43907) (Case: 00156347) (Case: 00149336) (Case: 00169221)
- Scheduled reports (Reports > Create Report > Scheduler) with a defined output of PDF are now emailed to recipients in the correct format. (JIRA ID: EM-39940) (Case: 00115228)
- The Interface Usage Report (Reports > Run Reports > Network Interfaces > Interface Usage) now displays accurate data for % utilization. (Case: 00156347) (Case: 00149336) (Case: 00169221) (JIRA ID: EM-43907)
- Addressed an issue where the availability data did not display in the Device at a Glance report. (Case: 00158556) (JIRA ID: EM-43222)

## Scheduler

- Addressed an issue where schedules configured to start before and run through the start of Daylight Savings Time, or schedules that started during Daylight Savings Time and ran through its exit, would not be properly adjusted after the Daylight Savings Time transition. (Case: 00137287, 00156979, 00157667) (JIRA ID:

EM-40671)

- Addressed an issue where a device that was put into maintenance, using a schedule with a patch window configured, remained in maintenance beyond its scheduled end time if the patch window never opened from an event of proper severity. (Case: 00112105) (JIRA ID: EM-38763)
- Addressed an issue where a device that was put into maintenance, using a schedule with a patch window configured, remained in maintenance beyond its scheduled end time if a reboot did not trigger the patch window. (Case: 00054618, 00057767, 00057975, 00088713, 00096437) (JIRA ID: EM-33397)
- Addressed an issue where the End Date was missing in some views of schedules that were configured as recurring schedules. (Case: 00097801) (JIRA ID: EM-37335)
- Addressed an issue where devices that had a maintenance schedule configured in multiple, overlapping schedules came out of maintenance early, as soon as the end time of the first schedule occurred instead of remaining in maintenance. (Case: 00079765) (JIRA ID: EM-36988)
- Addressed an issue where schedules configured in SL1 to be recurring and with no end date had their maintenance window end earlier than normal when another scheduling activity took place. (Case: 00105895) (JIRA ID: EM-36896)
- Addressed an issue where the calendars for Start Time and End Time for a schedule in the Scheduler Editor modal did not have vertical scrolling available for ease of viewing. A scrollbar was added to aid in viewing and using the calendars. (Case: 00090319, 00148108) (JIRA ID: EM-36325)
- Addressed an issue that occurred when devices were put into maintenance mode using a schedule with a patch window, and those devices remained in maintenance beyond the end of the scheduled exit time (if a reboot did not trigger the patch window). (Case: 00054618, 00057767, 00057975, 00088713, 00096437) (JIRA ID: EM-33397)
- Addressed an issue where, if you deleted a schedule using the SL1 API, SL1 did not fully remove the schedule from the database, as its associated task still existed, triggering another run of the schedule. With this release, the behavior was updated to ensure that a schedule deleted with the SL1 API is fully removed from the database. (Case 00045749) (Support ID: 180625) (JIRA ID: EM-32595)
- Addressed an issue where device maintenance schedules that were scheduled for a full day (24 hours/1440 minutes) would have their duration extended by a minute into the next day, causing the rest of the next day to be skipped. (Case: 00004075, 00072168) (Support ID: 170559) (JIRA ID: EM-26451)
- Modifying a non-active device maintenance schedule will no longer cause associated devices to go into maintenance. (Case: 00171348) (JIRA ID: EM-44207)

## ScienceLogic API

- Addressed an issue where PowerPacks that contained widgets could not be deleted using the API Browser. (Case: 00007118) (JIRA ID: EM-33127)
- API Addressed an issue where cleared events always show the user\_del attribute as null when those events are requested using the API. (Case: 00012656) (Support ID: 173936) (JIRA ID: EM-27968)
- The ScienceLogic API can now successfully fetch information from the Ticket Notes resource when the ticket note includes the special character for group separator. (JIRA ID: EM-19096) (Support ID: 127521)
- The ScienceLogic API now correctly creates and updates the credential in the /device/<id>/aligned\_app resource when the GUID is used as the <id>. (JIRA ID: EM-40079) (Case: 00127838)

- The ScienceLogic API now correctly creates and updates the credential in the /device/id/aligned\_app resource when the GUID is used as the ID. (Case: 00127838) (JIRA ID: EM-40079)

## Search

- SL1 now retains the search criteria on pages in the new user interface ("AP2") after you refresh the browser or navigate to another page and then return to the search page using the Back button in the browser. (Case: 00091267) (JIRA ID: EM-36526)

## Security

- Improved security for backup logs. (Case: 00171431) (JIRA ID: EM-44293)
- The password that you set during installation for SL1 on a MUD system no longer ages out immediately after installation completes. (Case ID: 00177639) (JIRA ID: EM-44435)

## SL1 Extended

- On SL1 Extended systems using Collector Pipeline, the last\_poll is now accurate and the Device Investigator now displays the correct date and time as the collection time. (Case: 00191613) (JIRA ID: CPL-567)

## SNMP Walker Tool

- Addressed an issue where the SNMP Walker tool generated the following error on devices using SNMP version 3 and empty Security Passphrase: /usr/bin/snmpwalk: (The supplied password length is too short.) Error generating a key (Ku) from the supplied authentication pass phrase. (Case: 00158634) (JIRA ID: EM-43263)
- Removed the License Count column from the "Subscription License Usage by Device" report because the column was not displaying the correct data. (JIRA ID: EM-40152) (Case: 00124338)

## System Update

- The patch import process now checks for the number of files associated with the patch file version ID, and will fail the import state if there are no packages associated with the version ID. This will prevent staging errors, because the staging button will not be enabled.
- The schema and data update scripts for post-deployment will now retry twice after failing initially with some delay in between retries. This will help prevent failed deployments and the need to subsequently attempt the deployment operation from the UI.
- System Update no longer includes code that rebuilds device relationships. System Update can now successfully perform updates on SL1 systems that include multiple component maps. System Update no longer fails when trying to rebuild device relationships and no longer consumes all disk space on the /tmp partition. (Case: 00090042) (JIRA ID: EM-36368)
- If there are no packages associated with a release (em7 and em7-os), staging and deployment complete successfully, and SL1 displays a message stating that there are no packages for SL1 version. (Case: 00135154) (JIRA ID: EM-40601) (JIRA ID: EM-40537)

## Tickets

- If you use the ScienceLogic API to update the `date_edit` field, you can now successfully set `"date_edit = now"`. (Case: 00100032) (JIRA ID: EM-37706)
- Addressed an issue where default ticket notifications continued to be sent after disabling Automatic Ticketing Emails on the Behavior Settings page. (Support ID: 141102) (JIRA ID: EM-22106)
- Addressed an issue where if you added a cloaked or uncloaked note to a ticket using the API, SL1 did not send any notifications for the ticket watchers. (Support ID: 113936) (JIRA ID: EM-16663)

## User Interface

- The PHP Developer Logs page (System > Tools > PHP Developer Logs) is now hidden from non-admin users. (JIRA ID: EM-44061) (Case: 00171630)
- References to the deprecated "EM7 Windows WMI Agent" have been removed from the classic user interface. (JIRA ID: EM-43253)
- The Navigation Tab Editor page (System > Customize > Navigation) now loads successfully with no error messages. (JIRA ID: EM-41213) (JIRA ID: EM-18921) (Support ID: 11998) (Support ID: 166951) (Case: 00003341) (Case: 00140995) (Case: 00004636) (Case: 00003267) (Case: 00098085) (Case: 00098098) (Case: 00098102) (Case: 00098196) (Case: 00098995) (Case: 00099239) (Case: 00099211) (Case: 00140995)
- In the new UI (AP2), when you navigate to a page that uses iframes, the page no longer defaults to the Device Manager page (Registry > Devices > Device Manager in the classic UI). (Case: 00164373) (JIRA ID: EM-43361)

## Windows Services

- Addressed an issue that was causing Windows services to still display on the Windows Services page even after the associated applications were uninstalled from Windows servers. (Case: 00067752) (JIRA ID: EM-34677)

---

## SL1 Extended Architecture

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

- **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 (data pull, config push, streamer, and communication via encrypted SQL) in SL1. 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.
- **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.

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

- **Compute Cluster.** Compute nodes are the SL1 appliances 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. The following services and features require the compute function:
- **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.

---

## PowerPacks in 11.1.0

Before upgrading to 11.1.0, please verify whether any PowerPacks currently running on your system are “newer” than the PowerPacks included in this SL1 update. If the PowerPack on your system is “newer” than the one included with the SL1 update, 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 PP you do not want to update and select its checkbox.
3. Click the **Select Action** field and choose *Change Collection State: Disabled (recursive)*. 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. In the Select Actions drop-down list, select Change Collection State: Enabled (recursive).
9. Click the **[Go]** button.

## New and Updated PowerPacks

The 11.1.0 release **includes** the following PowerPacks that are new or updated and included with the release:

- Cisco Base Pack, v213
- Cisco: Telepresence Endpoint v101
- Cisco: UC VOS v109
- Cisco: Wireless v103
- Host Resources v107
- Cisco: ACI v111
- Cisco: Base Pack v213
- Cisco: Cloud Services Platform (formerly Cisco: CSP-2100) v106
- Cisco: Wireless v103
- Cisco: TelePresence Endpoint v101
- Cisco: UC VOS v109
- Host Resource Core Pack v107
- Microsoft: Windows Server v112
- Service Level Management PowerPack v101, Revision 34
- VMware: vSphere v301

## Deprecated PowerPacks

The 11.1.0 release **deprecates** the following PowerPacks and removes 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 v112 PowerPack.

## PowerPacks Removed from the ISO

The 11.1.0 release **removes the following PowerPacks from the ISO**. These PowerPacks are still available for download from the customer portal:

- Cisco: ACI Dashboards
- Cisco: ACI Reports
- Cisco: CUCM Cisco Unified Communications Manager
- Cisco: TelePresence: Traps
- Cisco: IPSLA
- Cisco: Meeting Server
- Cisco: UCS Standalone Rack Server
- Dell EMC xTremIO
- LayerX Cisco CDR
- NetApp Base Pack
- Nutanix Base Pack
- PureStorage Flash Array

## Community PowerPacks Removed from the ISO

The 11.1.0 release **removes the following community PowerPacks from the ISO**:

- 3Com Device Classes - Base Pack
- Alcatel-Lucent Device Classes - Base Pack
- Alteon Monitoring - Base Pack
- APC Base Pack
- Aruba Monitoring - Base Pack
- AskEM7 Query Widgets
- Attachmate Device Classes - Base Pack
- Avaya Base Pack
- Avocent ACS Pack
- Avocent Monitoring - Base Pack
- BlueCat Monitoring - Base Pack
- Blue Coat Monitoring - Base Pack
- Brocade Base Pack
- Cisco VPN Monitoring

- Citrix Monitoring - Base Pack
- Coyote Point Monitoring - Base Pack
- Danaher Device Classes - Base Pack
- DEC Device Classes - Base Pack
- Dell OM Base Pack
- Dell OpenManage Legacy Monitoring - Old Base Pack
- Dell PowerConnect Base Pack
- Dell PowerVault Event Policies - Base Pack
- D-Link Device Classes - Base Pack
- EM7 Nmap Device Classes
- EM7 Virtual Device Classes
- Empire Device Classes - Base Pack
- Enterasys Device Classes - Base Pack
- Extreme Base Pack
- Fluke Networks
- Force 10 Monitoring - Base Pack
- Fortinet Base Pack
- Foundry Base Pack
- H3C Monitoring - Base Pack
- Hitachi Base Pack
- HP Base Pack
- HP-ISM Monitoring - Base Pack
- HP-UX Base Pack
- IBM Director Base Pack
- Intel Base Pack
- Konica Minolta Base Pack
- LANCOM Systems Device Classes - Base Pack
- Lannair Device Classes - Base Pack
- Lantronix Device Classes - Base Pack
- Liebert Monitoring - Base Pack
- Linksys Device Classes - Base Pack
- McAfee Monitoring - Base Pack
- MIB-2 Base Pack
- Motorola Device Classes - Base Pack
- NetBotz Base Pack

- NetScout Systems Device Classes - Base Pack
- Netscreen Base Pack
- Nokia Base Pack
- Printer Base Pack
- Redis: Base Pack
- Riverbed Monitoring - Base Pack
- ScienceLogic EM7 Base Pack
- SNMP Research Base Pack
- Tomcat Monitoring
- UCD-SNMP Base Pack
- VMware: vSphere Reports
- Vyatta
- Xerox Base Pack

## PowerPacks Removed from ISO Due to New Pricing

For 10.2.0 and later releases, only "Base" PowerPacks will be included with the platform ISO. The following PowerPacks have been removed from the 10.2.0 and later ISOs to comply with the new SL1 pricing model:

**NOTE:** If you are upgrading from a previous version of SL1, the 11.1.0 upgrade will not remove any existing PowerPacks.

- Aruba Base Pack
- Avocent ACS Pack
- BlueCat Base Pack
- Cisco VPN Pack
- Cisco: AppDynamics
- Couchbase Base Pack
- Coyote Point Base Pack
- Dell OpenManage Old Base Pack
- H3C Base Pack
- IBM Director Base Pack
- LifeSize Endpoint
- Microsoft: Active Directory Server
- Microsoft: DHCP Server
- Microsoft: DNS Server
- Microsoft: Exchange Server

- Microsoft: Exchange Server 2010
- Microsoft: SharePoint Server
- Microsoft: SQL Server
- Microsoft: SQL Server Enhanced
- Tomcat

To upgrade your license and download PowerPacks, contact your Customer Success Manager.

Documentation and release notes for each PowerPack are available at the [PowerPacks Support](#) page.

---

## Disabling the Knowledge Base

The Knowledge Base includes known security vulnerabilities. ScienceLogic no longer supports the Knowledge Base.

- If your first installation of SL1 was 8.9.1 or earlier, ScienceLogic strongly recommends that you disable the Knowledge Base. SL1 provides a setting in the `silos.conf` file to disable the Knowledge Base.
- For newer installations where the first installation was 8.9.2 or later, the Knowledge Base will be disabled by default.

**WARNING:** The Knowledge Base includes known vulnerabilities for cross-site scripting and SQL injection. ScienceLogic strongly recommends that you disable the Knowledge Base.

To disable the Knowledge Base:

1. Use SSH to connect to the Administration Portal and Database Server or All-In-One (all SL1 appliances that provide a web interface).
2. Use an editor like `vi` and edit the file `/etc/silos.conf`. In the LOCAL section, add the line:  

```
kbasedisabled=1
```
3. Use an editor like `vi` and edit the file `/etc/siteconfig/silos.conf.siteconfig`. In the LOCAL section, add the line:  

```
kbasedisabled=1
```
4. Open a browser session and log in to SL1.
5. From the hamburger menu (☰) in the upper right, select *Clear SL1 System Cache*.
6. Upon your next login, the Knowledge Base tab will not appear. Attempts to access the tab will result in an "Access Denied" error message.

---

## Upgrade Process for Systems Running 8.1.0 and Earlier

**WARNING:** ScienceLogic strongly suggest you contact Customer Support or your Customer Success Manager to plan your migration from CentOS (versions of SL1 prior to 8.1.1) to 11.1.0.

The 8.1.1 release included a complete update of the ScienceLogic appliance operating system from CentOS 5.11 to Oracle Linux. Major operating system components, including the database, web server, and High Availability/Disaster Recovery packages have been updated or replaced by new, industry-standard packages.

When upgrading from a version prior to 8.1.1, each appliance must be migrated to 8.9.0 and the Oracle Linux 7.5 operating system.

---

## Upgrade Process for Systems Running 8.1.1 and Later

**TIP:** For detailed instructions on planning an upgrade, best practices for upgrades, and executing an upgrade, see the chapter on *Upgrading SL1* in the *System Administration* manual or view that chapter [online](#).

If you are running 8.4.0 or earlier and require access to all ticket notes immediately after upgrading, contact ScienceLogic Customer Support for details on manually updating the database schema **before you upgrade**.

If you are running 8.4.0 or earlier and have added one or more custom firewall rules, such as a non-standard port for Phone Home Collectors, you must migrate these rules to firewalld **before you upgrade**. Please contact ScienceLogic Support for more information.

If you are upgrading from a version of SL1 prior to 8.6.0, you will have to import, stage, run the pre-upgrade script, and deploy the update twice: once to upgrade to 8.6.0 and then again to use a delta-less upgrade to the latest update release.

## Downloading SL1 Updates on SL1 Systems running 8.1.x - 8.5.x

To download updates for previous SL1 software 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.

You must upgrade your system to 8.6.0 and then upgrade again with the newer deltaless upgrade process.

Store the update files in a location that you can use to upload files to the SL1 system.

**NOTE:** These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least 3 days before upgrading.

**TIP:** For detailed instructions on planning an upgrade, best practices for upgrades, and executing an upgrade, see the chapter on **Upgrading SL1** in the **System Administration** manual or view that chapter [online](#).

## Downloading SL1 Updates on SL1 Systems Running 8.6.0 or Later

If your SL1 System is running version 8.6.0 or later, you can download a single update file and update your SL1 system to the latest release.

Before you can load a patch or update onto your instance of SL1, you must first download the patch or update to your local computer:

1. Log in to the ScienceLogic Support site at <https://support.sciencelogic.com/s/>. Use your ScienceLogic customer account and password to access this site.
2. Select the **[Product Downloads]** button, select the **Product Downloads** menu, and choose *Platform*.
3. Find the release you want and click its name.
4. In the **Release Version** article, click on the link for the release image or release patch you want to download. Scroll to the bottom of the page.
5. Under **Files**, click the link for the file you want to download. The file is downloaded to your local computer.

**NOTE:** These steps do not affect the performance of SL1. ScienceLogic recommends that you perform these steps at least three days before upgrading.

**TIP:** For detailed instructions on planning an upgrade, best practices for upgrades, and executing an upgrade, see the chapter on "Upgrading SL1" in the **System Administration** manual or view that chapter [online](#).

---

## Recently Deprecated Features

### 8.14.0

- Deprecated the SNMP-based version of the ScienceLogic Support PowerPack (EM-30510)
- The SSH Tool has been removed from the Device Toolbox (Registry > Devices > Device Manager > wrench icon > Toolbox). (Case 00022135) (Support ID: 176020), (EM-29178)
- The Content Management page appears in the user interface but has been deprecated. Updates to the user interface are now included in platform updates.

## 10.1.0

- The Content Management page no longer appears in the user interface.
- Deprecated harProviderSearch and deviceSearch and replaced with override search. (SLUI-7404)
- The Video Reports PowerPack is no longer included with ISO builds. (SOL-6778)
- The Devices > Agent tab is now part of Device Settings (SLUI-6386)

## 11.1.0

- Removed flashed-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 deprecates the following PowerPacks and removes them it 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 v112 PowerPack.

© 2003 - 2023, ScienceLogic, Inc.

All rights reserved.

#### LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

#### Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

#### Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: [legal@sciencelogic.com](mailto:legal@sciencelogic.com). For more information, see <https://sciencelogic.com/company/legal>.





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

International: +1-703-354-1010