



SL1 10.2.0 Release Notes

SL1 version 10.2.0 , Revision 5

Table of Contents

Upgrade Notes	4
Caveats	6
Known Issue in 10.2.0	7
Critical Known Issue with the Gen 1 Agent	7
Major Known Issue with Data Collection	8
Disable Concurrent SSH Collection	8
Manually Prune the files in /var	9
New Features in 10.2.0	9
Action Runner (Run Book Automation)	9
Admin Notifier	9
Agent and Agent Pipeline	9
Anomaly Detection	10
Appliance Manager	10
Backups	10
Business Services	10
Collection	11
Collector Pipeline	11
Component Device Relationships (DCMR)	11
Concurrent SNMP Collection	11
Concurrent PowerShell Collection	12
Credentials	12
Dashboards	12
Deployment	13
Developer Logs	13
Device Classes	13
Device Investigator	13
Device Management	14
Disaster Recovery and High Availability	14
Discovery	15
Events	15
Extended Architecture	15
Global Manager	15
IP Address	16
Logging	17
Maps	17
MariaDB	17
MySQL	17
New UI	17
Platform	17
PowerPacks	19
Publisher	19
Reports	20
Security	20
SNMP Traps	20
Subscription Billing	20
System Update	20
User Interface	21
Issues Addressed in 10.2.0	21
Access Hooks and Access Keys	21
Agent	21

Alerts	21
Backups	21
Business Services	22
Collection	22
Collector Pipeline	22
Credentials	22
Dashboards	23
Database Server	23
Devices	23
Dynamic Applications	23
Email	23
Events	23
High Availability and Disaster Recovery	24
IP Address	24
Maps	24
MariaDB	24
MIB Compiler	24
PhoneHome Collector	25
Platform	25
PowerPacks	26
Process Monitoring	26
Reports	26
REST API	27
Security	27
Ticketing	27
Topology Maps	27
User Policies	27
Windows	27
User Interface	28
SL1 Extended Architecture	28
PowerPacks in 10.2.0	29
Disabling the Knowledge Base	31
Upgrade Process for Systems Running 8.1.0 and Earlier	32
Upgrade Process for Systems Running 8.1.1 and Later	32
Downloading SL1 Updates on SL1 Systems running 8.1.x - 8.5.x	32
Downloading SL1 Updates on SL1 Systems Running 8.6.0 or Later	33
Recently Deprecated Features	34
8.9.0	34
8.9.2	34
8.10.0	35
8.12.0	35
8.14.0	35
10.1.0	35
Known Issues	35

Overview

This document includes the following topics:

Upgrade Notes	4
Caveats	6
Known Issue in 10.2.0	7
New Features in 10.2.0	9
Issues Addressed in 10.2.0	21
User Interface	28
SL1 Extended Architecture	28
PowerPacks in 10.2.0	29
Disabling the Knowledge Base	31
Upgrade Process for Systems Running 8.1.0 and Earlier	32
Upgrade Process for Systems Running 8.1.1 and Later	32
Recently Deprecated Features	34
Known Issues	35

TIP: 10.2.0 includes delta-less updates. If you have already deployed 8.6.0 or later, you can import, stage, and deploy a single update file to update to 10.2.0.

Upgrade Notes

- **If you are running a version prior to 8.12.0, do not install 10.2.0 if you do not plan to immediately consume 10.2.0.** After you import the 10.2.0 release, all appliances in your SL1 system will now use the **new system update**. After you import the 10.2.0 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 10.2.0 and its new System Updates tool, you must have already imported, staged, and deployed 8.6.0 or a later release.**
- You can install 10.2.0 in AWS RDS/Aurora environments.
 - You can upgrade from 10.1.x running in AWS RDS/Aurora environments to 10.2.0 running in AWS RDS/Aurora environments.
 - However, you cannot upgrade from 8.14.x running in AWS RDS/Aurora environments directly to 10.2.0 running in AWS RDS/Aurora environments.

- 10.2.0 includes important security updates. **You must reboot all AIO appliances and all appliances in the SL1 Distributed Architecture after deploying 10.2.0.** 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/0TO0z000000E6w7GAC/known-issues) for SL1 (<https://support.sciencelogic.com/s/topic/0TO0z000000E6w7GAC/known-issues>) before installing a new update. For Known Issues specific to 10.2.0, 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?	Requires MariaDB Upgrade?
8.10.0	Yes	No
8.10.1	Yes	Yes
8.10.2	No	No
8.10.3	No	No
8.12.0	Yes	Yes
8.12.1	Yes	No
8.12.2	No	No
8.14.0	Yes	Yes
8.14.1	No	No
8.14.2	No	No
8.14.3	Yes	No
8.14.4	No	No
8.14.5	No	No
8.14.6	Yes	No
8.14.7	No	No
8.14.8	Yes	Yes
8.14.9	Yes	No
10.1.0	Yes	Yes
10.1.1	Yes	No
10.1.2	No	No
10.1.3	Yes	No
10.1.4	Yes	No
10.1.5	Yes	No
10.2.0	Yes	Yes

Caveats

Consider the following caveats before deploying 10.2.0:

- 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.
- 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/1161> 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 10.2.0 .**
- 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](#).**
- Global Manager, SL1 Extended Architecture, the new UI, and PowerFlow do 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 Issue in 10.2.0

Critical Known Issue with the Gen 1 Agent

Upgrading to the 10.2.x causes a service outage with Gen 1 Agents. After upgrade to the 10.2.x line, SL1 can no longer communicate with Gen 1 Agents. ScienceLogic recommends that customers using the Gen 1 Agent not upgrade to the 10.2.x line.

For details on about the known issue and status updates on the fix, please contact ScienceLogic Customer Support.

Major Known Issue with Data Collection

10.2.0 includes a major Known Issue that can cause performance degradation and a potential outage for Data Collectors. (Case: 00170950)

The Concurrent SSH Collection feature is enabled by default on 10.2.0 systems, both upgrades and new installations. The Concurrent SSH Collection feature generates temporary files for each connection attempt to the monitored devices. These temporary files can eventually use all the space in the /var directory of each Data Collector, which can cause performance degradation and a potential outage.

Currently, only the Linux Base Pack PowerPack v103 and later uses concurrent SSH collection.


For full details, see the Knowledge Base Article: <https://support.sciencelogic.com/s/article/5498>

To prevent performance degradation and a potential outage, perform the following workaround steps:

- **Disable Concurrent SSH Collection.** This does not affect PowerPacks that use SSH. When concurrent SSH collection is disabled, the Linux Base Pack PowerPack will use standard SSH collection.
- **Manually prune the temporary files on each Data Collector.** Currently, only the Linux Base Pack PowerPack uses concurrent SSH collection. If you know which Data Collectors are performing collection for the Linux Base Pack PowerPack, you can perform these steps on those Data Collectors. If you do not know which Data Collectors are performing collection for the Linux Base Pack PowerPack, you can safely perform these steps on all Data Collectors.

Disable Concurrent SSH Collection

To disable Concurrent SSH Collection:

1. Log in to the Administration Portal.
2. Go to the Process Manager page (System > Settings > Processes).
3. Find the process "Data Collection: SSH Collector" and select its wrench icon ().
4. Set the **Operating State** to *Disabled*.
5. Save your change.
6. Wait 10 minutes for all collection processes to complete

CAUTION: : If you are monitoring a large number of Linux devices, for example, 100 devices or more, you must re-balance the Data Collectors after stopping Concurrent SSH Collector. For details, see the chapter on *Collector Groups and Load Balancing* in the **Systems Administration** manual.

Manually Prune the files in /var

Currently, only the Linux Base Pack PowerPack v103 and later uses concurrent SSH collection. If you know which Data Collectors are performing collection for the Linux Base Pack PowerPack, you can perform these steps on those Data Collectors. If you do not know which Data Collectors are performing collection for the Linux Base Pack PowerPack, you can safely perform these steps on all Data Collectors.

To manually prune the temporary files in the /var directory on each Data Collector:

1. Go to the console of the Data Collector or use SSH to access the Data Collector.
2. To delete all the temporary files in the Docker container for Concurrent SSH Collection, enter the following at the shell prompt:

```
for container in $(sudo docker ps -a | grep silo_ssh_collector | awk '{ print $1}'); do sudo docker rm -f $container; done
```

3. Notice that /var has more free space now.
4. Perform these steps on either each Data Collector performing collection for the Linux Base Pack PowerPack or on each Data Collectors in your SL1 system.

New Features in 10.2.0

Action Runner (Run Book Automation)

- Added a new "Custom Links" section in the Action Runner that appears on the Events page, the Devices page, and the Activity Center. With this update, an administrator can define URLs to external systems that will appear in this new "Custom Links" section. When a user clicks one of these links, it opens the URL in a new browser window.

Admin Notifier

- The new Admin Notifier alerts administrator users upon login to any issues on monitored SL1 systems that could lead to an outage (for example, if the database is running out of space). After login, the Admin Notifier displays a banner if any of the filtered events are found so that you can take action on the issues immediately."

Agent and Agent Pipeline

- Added a new option in the Agents section of the Device Investigator Settings tab: Upload Interval. Specify how often the agent should upload data. You can upload a data snapshot every 20 seconds, or upload a data summary every 60 seconds. This option uses an improved data format that requires less resources. The agent continues to internally collect and poll data every 20 seconds, but the agent summarizes and uploads that data every 60 seconds. There is no data loss even though the data is uploaded less frequently. Note that the Upload Interval option is currently available only with version 132 or later of the Windows agent.
- The SL1 Windows agent and SL1 Linux agent can upload polled data output and log data output in a compressed format, reducing the amount of data that each agent uploads.

- The diagnostic command for the Windows agent and the diagnostic option for the Linux agent were updated to include the current polled data configuration from the agent. This configuration information can help troubleshoot issues with local execution of Dynamic Applications.
- The Windows 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.
- Updated the Windows agent to load the shared library, `sl_intercept.dll`, into processes on all versions and builds of Windows 10.
- When you uninstall a Windows agent, SL1 removes the `C:\program files\sciencelogic\siloagent` directory during uninstall. Previously this directory was left on the agent, which sometimes created duplicate devices if you wanted to re-install the agent on the same device.
- When you uninstall a Linux agent, SL1 removes the `/etc/scilog/` and `/opt/scilog/` directories during uninstall. Previously these directories were left on the agent, which sometimes created duplicate devices if you tried to re-install the agent on the same device.

Anomaly Detection

- Updated the anomaly detection model selection process to further reduce the detection of false-positive anomalous behavior.
- Added a timespan selector to the machine learning line graph that appears on the Machine Learning tab of the Device Investigator and the Anomalies tab of the Service Investigator. With this update, you can choose from the following timespan options for the machine learning graph: 1 hour, 3 hours, 6 hours, 12 hours, 24 hours, and 2 days. The default value is 24 hours.
- Added the ability to filter devices on the Machine Learning page based on whether anomaly detection is enabled or disabled.

Appliance Manager

- The Appliance Manager page (System > Settings > Appliances) highlights SL1 appliances that are running a different version of SL1 than the Database Server.
- When Apple employees die, do their lives HTML5 in front of their eyes?

Backups

- SL1 now supports full backups and configuration backups on Database Servers configured for Disaster Recovery.
- User can specify the number of backups to retain before new backups overwrite the older backups.
- Configuration backups now include Business Service information.

Business Services

- Added the ability to create custom service models from the Business Services page, to enable users to create service hierarchies that accurately reflect the service structures within their organization. With this update, users can use a wizard to build service models with multiple nested or connected service levels, each of

which users can label according to their needs. As with other service types, users can create templates from existing service models and then use those templates to create new service models. For more information about creating custom service models, see the Monitoring Business Services manual.

- Added a Map view on the Overview page in the Business Service Investigator. With this update, users can now use a drop-down menu on the Overview page to view their service structure as a hierarchical map rather than the default sunburst view.
- The Business Services PowerPack was updated to meet current ScienceLogic engineering standards.
- Added a Custom Attributes tab to the Service Investigator. This tab displays a list of the custom attributes that are aligned with your service. On this tab, you can align additional custom attributes to the service, edit the custom attributes' values, and unalign some custom attributes from the service. You can use custom attributes when importing services from an integrated system to handle incoming properties that are not defined in SL1. Only those custom attributes with a Resource Type of "Service" can be aligned to services.

Collection

- The SL1 Extended Architecture can now publish alerts, CPU, memory, and swap data collected by Dynamic Applications. The published data can be consumed by AIML and the new Event Engine.
- Added support for Dynamic Component Mapping during discovery with Concurrent SNMP or Concurrent PowerShell.
- In SL1 Extended Architecture, users can toggle on or toggle off the publishing of Dynamic Application data for CPU, Dynamic Application data for memory, Dynamic Application data for SWAP. If AIML is enabled, SL1 automatically publishes performance data from Dynamic Applications. If the new Event Engine is enabled, SL1 automatically publishes alert data from Dynamic Applications.

Collector Pipeline

- The SL1 Extended Architecture includes update and fixes to ``interface-alert``, ``interface-store``, ``interface_magic``, and ``streamer-push`` services.
- SL1 Extended architecture now supports alerting from availability policies and network interface policies.
- To conserve memory, Collector Pipeline now uses a new rollup service (`sl_rollup_utils`) to aggregate hourly and daily rollups.

Component Device Relationships (DCMR)

- SL1 includes a new setting in the Behavior Settings page (System > Settings > Behavior). The **Component Device Map Update Mode** field allows users to specify the method in which SL1 updates relationship maps. The default method, using DCM-R Triggers, is the same method used in previous releases. The new Periodic update method is designed for customers experiencing data base performance issues related to large number of changes in device topology.

Concurrent SNMP Collection

- Added a new feature, Concurrent SNMP Collection. This type of SNMP collection allows multiple collection tasks to run concurrently with a reduced load on Data Collectors. This feature is disabled by default.

- When Concurrent SNMP Collection is enabled, each enabled Data Collector now includes four instances of the SNMP Collector container, to enable parallel processing. A single failed task will no longer prevent other tasks from completing, and throughput for SNMP collection is greatly improved.
- Added support for Dynamic Component Mapping during discovery with Concurrent SNMP or Concurrent PowerShell.
- Improvement to shrink the size of one of the containers that performs concurrent Dynamic Application collection and can run on each Data Collector.
- Added a new logging library, `sl_logging` that allows parallel logging for Concurrent SNMP collection. Logs will continue to appear in the System Logs page (System > Monitor > System Logs) and will continue to trigger internal alerts as necessary.

Concurrent PowerShell Collection

- Added support for Dynamic Component Mapping during discovery with Concurrent SNMP or Concurrent PowerShell.

Credentials

- Added five new "universal" credential types for monitoring Aliyun/Alibaba Cloud, AWS, Microsoft Azure, Citrix XenServer/XenCenter, and IBM Cloud/SoftLayer resources. These credential types, which can be used for each of the guided discovery workflows available in SL1, include fields and labels that are specific to each platform. For example, the AWS credential includes fields such as "AWS Access Key ID" and "AWS Secret Access Key" rather than the generic SOAP/XML credential fields "HTTP Auth User" and "HTTP Auth Password".
- Redesigned the Credential Editor page layout for each credential type. As part of this redesign, all fields that are common across credential types now appear at the top of the page, and a Credential Tester panel is now embedded on the right side of the page.
- When creating an SSH/Key credential in the Next UI, SL1 will validate that the private key is entered in the correct format. The user will be able to save the credential only if the private key is correctly formatted.

Dashboards

- Updated the timespan picker in dashboards to enable users to customize the timespan of data displayed in the dashboard. With this update, you can either specify a custom timespan, select an absolute timespan such as "Last Week" or "Today so far", or select a relative timespan such as "Last hour" or "Last 7 days".
- In dashboards, when creating a Services widget, the "Type" drop-down menu has been relabeled "Service Type", and you can now search for and select any labels associated with custom service models in addition to the existing Service Types. (SLUI-8212)
- In dashboards, you can now select a Rollup Frequency of "Auto" for time series metrics. When you do so, this option will summarize the data displayed to an appropriate level for the requested time range. For example, time ranges of 2 days or less will display raw data, time ranges of more than 2-45 days will display hourly data, and time ranges of more than 45 days will display daily data.
- In dashboards that use the Event Table widget, service events will now display "Service" in the Event Type column instead of "HarProvider".

- In dashboards, you can now select a specific device whose configuration information you want to appear in a Configuration Table widget. When you do so, the name of the device will appear in the title at the top of the widget. If you select more than one device, only the first device you select will appear in the widget.

Deployment

- For SL1 Extended Architecture, Updated IAC framework to support Bitnami endpoint due to vendor changes.
- Enabled SNMPD on the SL1 Management nodes, Compute nodes, and Storage nodes for self-monitoring.
- 10.2.0 includes an upgrade to Kubernetes 1.16.
- When updating an SL1 Extended Architecture, the Management Node automatically updates versions of internal packages (currently Ansible, Docker, docker-compose, Helm, helm-diff, k9s, kubectrl, RKE, Scylla) as necessary.

Developer Logs

- Added a new SL1 Developer Logs page that allows users to enable and download developer log files in the SL1 Next UI. You can access this page by going to System > Tools > SL1 Developer Logs on the Advanced Menu. For more information, see the System Administration manual.

Device Classes

- Added new default icons for some device classes. SL1 will use these default icons when more specific icons do not exist for those device classes.

Device Investigator

- The Device Investigator sidebar has been moved to the right side of the page and is now collapsible. This sidebar also now includes the timespan selection that was previously at the top of the page. To add, combine, or rearrange Device Investigator panels, click the Edit button on the right-hand sidebar to expand the configuration panel.
- The Device Investigator now includes a new overview panel. This panel appears at the top of the Device Investigator and includes basic information about the device, including its IP address, device class, organization, current health and collection statuses, and number of events by severity.
- The Device Investigator now includes a responsive column layout that can display multiple panels in side-by-side columns, rather than just a single column of panels, that will automatically adjust based on your browser width.
- Updated the Device Investigator to give users much more control over its layout. With this update, users can now determine the relative size (Small, Medium, or Large) of each panel in the Device Investigator, or collapse or expand individual panels as needed. In addition, the panel layout is now responsive, so columns are automatically resized whenever a metric panel is added or removed or whenever the panels are rearranged on the page.
- Added leaderboard chart panels for interfaces, file systems, and indexed Dynamic Applications in the Device Investigator. Some panels now include a header button that enables users to switch from a leaderboard bar chart to a line chart or vice versa.

- Added the ability to save a custom Device Investigator layout and apply it to one or more Device Categories, Device Classes, or devices. When you do so, you can give the layout a custom name and apply the layout either to just yourself or to all users with the appropriate access hooks aligned to their profiles. If you apply the layout only to yourself, the layout will be considered private; if you apply it to all users with the appropriate access hooks, it will be labeled as public. When no custom layout is specified, the system will use the default Device Investigator layout.
- Redesigned the Device Investigator by moving the metric panel list and the timespan filter to a collapsible sidebar drawer on the right-hand side of the page, so they take up less space on the page when not needed. This sidebar also includes a header that displays the alignment type (Default, Custom, Device, Class, or Category) and visibility settings (Default, Public, or Private) for the sidebar.
- Updated the timespan rollup frequency options for metric data that appears in the Device Investigator. Previously, metric panels displayed raw data, regardless of the selected timespan. With this update, unless users specify that they want to always display raw data, Device Investigators with a timespan of 2-45 days will display hourly data rollups and those with a timespan of longer than 45 days will display daily data rollups.
- Performance upgrades were made to the Device Investigator and other device-related pages to enable them to load faster.

Device Management

- Added controls to prevent an unauthorized user from deleting any device via the API.

Disaster Recovery and High Availability

- Added new features to SL1 systems configured for Disaster Recovery. The SL1 system displays the following information:
 - Cluster Maintenance is detected. Provides information to put nodes into maintenance and out of maintenance.
 - Reboot of Primary Database Server. Provides prompt to manually promote the primary Database Server and information on how to do so.
- Improvements to High Availability and Disaster Recovery systems. The file `/opt/em7/gui/ap/www/health/index.em7` now returns the hostname and specifies which Database Server is the primary.
- 10.2.0 includes an update to drbd-proxy.
- SL1 now supports full backups and configuration backups on Database Servers configured for Disaster Recovery.
- The ScienceLogic Support PowerPack generates an event if the configuration files differ on each Database Server in a cluster (High Availability or Disaster Recovery).
- Upon login, High Availability and Disaster Recovery systems display a "Message of the Day". This message can:
 - Display the current status of the system (for example, in maintenance, in split-brain)
 - Alerts about license expiry

- Remind users to manually promote a Database Server after restart
- Instruct users how to fix split-brain error

Discovery

- Updates to the new user interface so that discovery aligns to our new user interface patterns and leverage common new user interface components.
- Streamlined the guided discovery process by removing or consolidating steps, and added a header that displays the guided discovery steps and where the user is in the process.
- Updated the logo icons used for some of the default guided discovery workflows.
- The guided and unguided discovery workflows were updated to include a list of workflow steps at the top of each page. After completing a step, users can click on the step number to return to that step, but they cannot advance to steps they have not yet completed.
- Added a new "Guided Discovery Workflow Base Pack" PowerPack. This PowerPack includes several guided discovery workflows and all of their associated icons and is installed by default on your SL1 system.

Events

- SL1 Extended Architecture now supports alerts from data that uses Collector Pipeline and Agent Pipeline (Availability data, Network Interface data, and performance data from Dynamic Applications.) Alerts messages are forwarded to device logs and compared to event definitions, as they are for the SL1 Distributed Architecture. Event records are then stored in MariaDB, as they are for the SL1 Distributed Architecture.
- The SL1 Extended Architecture uses a new, scalable, kafka-based process to manage alert data from Dynamic Applications,
- The SL1 Extended Architecture can now publish alerts. The published data can be consumed by the new Event Engine.

Extended Architecture

- SL1 Extended Architecture now supports alerts from data that uses Collector Pipeline and Agent Pipeline (Availability data, Network Interface data, and performance data from Dynamic Applications.) Alerts messages are forwarded to device logs and compared to event definitions, as they are for the SL1 Distributed Architecture. Event records are then stored in MariaDB, as they are for the SL1 Distributed Architecture.

Global Manager

- SL1 Global Manager has been completely redesigned and updated for the SL1 Next UI. To enable this new version of Global manager, you must set the GLOBAL_MANAGER=enabled variable in the nextui.conf file on your SL1 system. The following bullets describe the changes made to Global Manager.
 - The Global Manager Stacks page has been completely redesigned and is now accessible by clicking the Global Manager Stacks icon at the top of the left navigation bar.

- This page displays a list of all of your Global Manager stacks and enables you to add new stacks or edit or delete existing stacks.
 - To view this page, users must have the GM_STACK_VIEW hook aligned to their user policies.
 - To add, edit, or delete stacks, they must have the GM_STACK_ADD, GM_STACK_EDIT, and GM_STACK_REMOVE access hooks aligned.
- To use the updated Global Manager, both the Global Manager appliance and all SL1 appliances in the managed stacks must be running 10.2.0.
- The "ScienceLogic: Global Manager" PowerPack has been updated to provide a second, simpler Global Manager setup and discovery process, in addition to the existing process. With this update, you can create Global Manager stacks immediately after SNMP discovery of All-in-One or CDB SL1 systems. The Run Book Actions in the PowerPack will automatically add entries to the Global Manager stacks.
- To determine the version of SL1 being used for Global Manager, go to the About page and look for the "appliances" value in the Platform section.
- For Global Manager systems, a "Global View" toggle button has been added to the top of the Devices, Events, and Dashboards pages. When toggled on, users can view the devices or events that exist across all Global Manager stacks or dashboards that have been created and saved in Global Manager mode; when toggled off, only the devices discovered on the Global Manager system itself, as well as events and dashboards for those devices, will display. When in Global Manager mode, device IDs will be preceded by the stack number. You cannot add devices while in Global Manager mode.
- In Global Manager mode, the Device Investigator, Settings, Configs, Events, Interfaces, and Software tabs display for globally managed devices. From these pages, if users want to edit the device, they can click a new "Manage" button, which will open the device in non-Global Manager mode in a separate window. After editing the device and saving their changes, user can either close the window to return to Global Manager mode or keep it open to continue in non-Global Manager mode.
- In Global Manager mode, stack details now appear in the Info drop-down box on the Device Investigator for all globally managed devices.
- In Global Manager mode, clicking on an event on the Events page will open the Event Investigator page. From this page, you can acknowledge or clear the event, add or edit notes about the event, and view any logs or vitals graphs related to the event or the device associated with it. When in Global Manager mode, the Tools panel and Actions button will not appear on the page.
- In Global Manager mode, you can create external tickets for events. When you do so, the external ticket is aligned to the event in non-Global Manager mode as well.
- In Global Manager mode, you can view or create dashboards that include data from all stacks. When creating a dashboard in Global Manager mode, you can create Device, Device Components, Event, Interface, and File System widgets for the dashboard.

IP Address

- After changing the IP address of a Data Collector or Message Collector in the Web Configuration Utility, the IP address is updated correctly in all required files and database tables.
- After changing the IP address of a Data Collector or Message Collector at the shell prompt, the IP address is updated correctly in all required files and database tables.

Logging

- To allow diagnostic access to logs from the nextui service, the nextui services now writes logs to both /var/log/em7/nextui.log and /var/log/messages. The log files in /var/log/em7/nextui.log are rotated.
- Added a new logging library, sl_logging that allows parallel logging for Concurrent SNMP collection. Logs will continue to appear in the System Logs page (System > Monitor > System Logs) and will continue to trigger internal alerts as necessary.

Maps

- Redesigned the Maps page with an updated and responsive toolbar, new icon buttons, an expanded search bar, and a new drop-down Selections panel.
- Redesigned the Selections menu in maps to use a new toolbar icon instead of a drop-down menu. Additionally, added a new Processes tab to the Selections menu to display a list of processes for selected devices.

MariaDB

- This release includes an upgrade to MariaDB. The upgrade is 10.4.18.

MySQL

- The mysql command and mysql_upgrade command no longer requires the user to input a username and password when connecting to the local database. (EM-38232)

New UI

- Updated the look and functionality for some drop-down lists to include a search field and multi-select checkboxes and to display selected fields at the top of the list.
- Updated the look and functionality for date and time selectors.
- Updated the look and functionality of the "share with" drop-down menus that are used to determine whether maps and dashboards are public, private, or shared with one or more specific organizations.
- Updated the look and functionality of the wizards used in SL1 to guide users through a process, such as guided discovery. (
- Added a new SL1 Developer Logs page that allows users to enable and download developer log files in the SL1 Next UI. You can access this page by going to System > Tools > SL1 Developer Logs on the Advanced Menu. For more information, see the System Administration manual.
- Added new default icons for some device classes. SL1 will use these default icons when more specific icons do not exist for those device classes.

Platform

- The process "EM7 Core: Daily Maintenance" includes tasks that prune unused data from the database. This process has been improved to run jobs in parallel to improve efficiency and prevent blocking of other jobs.

- SL1 now includes automated self-healing. Self-healing jobs run on Database Servers and All-in-One appliances and can :
 - automatically set "s-em7-core:s-em7-core" as the owner and group for the file "silo.log"
 - automatically restart the process em7_patch_manager if it is stuck in deactivating mode
 - automatically kill queries that have run for longer than 1 hour and logs each query
- The latest version of the SL Self-Monitoring PowerPack (ScienceLogic Support Pack)
 - Monitors unsent emails from SL1 . Every five minutes, SL1 checks the mail queue for outgoing email. If more than 100 email messages are queued and unsent, SL1 generates an event.
 - Monitors DNS entries for Database Servers and All-In-One appliances.
 - Triggers an event is an SL1 appliance in not currently monitored by SL1 .
 - Ensures that all SL1 appliances appear in the Device Manager page (Registry > Devices > Device Manager)
 - Generates alerts if any of the following files differ on an SL1 appliance::
 - /etc/my.cnf.d/silo_mysql.cnf
 - /etc/silo.conf
 - /etc/siteconfig/mysql.siteconfig
 - /etc/siteconfig/siloconf.siteconfig
 - Generates alerts if a query runs for longer than 15 minutes. Also logs the query .
 - Generates an event if the mariadb log includes error messages or warning messages.
 - Generates an event if the configuration files differ on each Database Server in a cluster (High Availability or Disaster Recovery).
 - Monitors the performance and configuration of the MariaDB database.
 - Generates a major event if CRM configuration or corosync configuration are out-of-date.
 - Ensures that all SL1 appliances appear in the Device Manager page (Registry > Devices > Device Manager).
 - Triggers an event is a process is killed by SL1 due to an out-of-memory condition.
 - Auto-clears events about licensing when a license is renewed.
 - Ensures that all SL1 appliances are running the same version of SL1 .
- Added a new feature to SL1 System Administration. Users can now change the IP address of a Data Collector or Message Collector in the Web Configuration tool, and the IP will then be updated in:
 - /etc/silo.conf
 - /etc/sysconfig/network-scripts/ifcfg-ens160
 - /etc/hosts
 - the ip column in the master.system_settings_licenses database table

And licenses will be unaffected.

- Added a new feature to SL1 System Administration. Users can now use SSH to connect to a Data Collector or Message Collector and change the IP address by entering:

```
update_IP <interface ID> < new IP address>
```

The IP address will then be updated in:

- /etc/silo.conf
 - /etc/sysconfig/network-scripts/ifcfg-ens
 - /etc/hosts
 - the ip column in the imaster.system_settings_licenses database table
- The Appliances Manager page (System > Settings > Appliances) includes a new column that specifies whether an SL1 appliance requires rebooting. This status is updated every 30 minutes. For SL1 appliances that require rebooting, the mouse-over test will display the list of new services or packages that require reboot, the current kernel version, and the date and time of the last reboot.
 - Eventing services introduced in this release for SL1 Extended architecture address an issue related to deadlocks caused by a process that updated Device Health. A new process, "Update Device State", has replaced the previous process. This process runs at a regular interval, which you can configure in System > Settings > Processes.

PowerPacks

- Added the ability to include Device Investigator layouts, universal credential fields, guided discovery workflows, and guided discovery workflow icons in a PowerPack. With this update, two new content options--"Credential Fields" and "Guided Discovery Workflow"--appear on the PowerPack Editor menu, where you can add those respective items. You can add Device Investigator layouts and guided discovery workflow icons by using the existing AP Content Objects content option on the PowerPack Editor menu. As with other content types, after you have added any of these new content options to a PowerPack, you can then export the PowerPack from one SL1 system and then import and install it and all of its included content onto another SL1 system.
- Added a new "Guided Discovery Workflow Base Pack" PowerPack. This PowerPack includes several guided discovery workflows and all of their associated icons, and is installed by default on your SL1 system.

Publisher

- The latest version of Publisher includes Kubernetes Operators. Subscription management has been redesigned. Supported egress target remains Kafka-only. To deserialize SL1 messages from their designated Kafka Topic(s) still requires the publisher client library and customer development.
- Messages sent from SL1 Publisher are encoded with the sl_schema_registry library. To read the messages that are sent to a Kafka topic, you must install the sl_schema_registry library and a Kafka Python library. You can now download the sl_schema_registry library and its documentation as a compressed file (sl_schema_registry_whl_and_docs.zip) from the ScienceLogic Support Site.
- SL1 Publisher is now included in the standard SL1 platform. Previously, this functionality required a separate package installation.

Reports

- SL1 10.2.0 build contains the "EM7 Performance Report PowerPack" version 101. EM-38816

Security

- In the **Behavior Settings** page (System > Settings > Behavior) removed the option "MD5" from the field **Password Hash Method**. **Password Hash Method** now defaults to "SHA-512". Existing passwords will remain in the hash type in which they were generated until the password is changed whereby the current setting will be used.

SNMP Traps

- Added a new feature to SL1 for SNMPv3 traps. After a user selects the SNMP V3 Trap Config Push button in the Credential Management page (System > Credentials page > Actions 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.

Subscription Billing

- A new page for viewing and managing subscription license usage has been added under Manage > Subscription Usage. This page shows license usage over time and the current license assigned to each device in SL1.
- The payload of data sent to ScienceLogic by the subscription usage process now includes the earliest license expiry date for all appliances in the SL1 system.
- Various performance improvements have been made to the subscription usage process.
- The payload of data sent to ScienceLogic by the subscription usage process now includes the latest output of the system status script from all appliances in the SL1 system.

System Update

- Updated MariaDB upgrade script to improve the SL1 upgrade process and timing.
- Improved the system update process to run more efficiently.
- Improved the Import process.
- To conserve disk space, SL1 will automatically prune Docker container images after System Update has completed.
- To conserve disk space, update files are automatically removed from the /tmp directory of SL1 appliances after the import process begins.
- The Appliances Manager page (System > Settings > Appliances) highlights SL1 appliances that are running a different version of SL1 than the Database Server.
- The Appliances Manager page (System > Settings > Appliances) includes two new columns: Needs Reboot and Task Manager Paused?

- The script "module_upgrade_mariadb" that performs an automatic upgrade of MariaDB on all relevant SL1 appliances now includes multiple improvements for 10.2.0.
- A new script, monitor_upgrade_mariadb, allows users to monitor the status of the MariaDB upgrade.
- SL1 includes a troubleshooting script for System Update.

User Interface

- The End User License Agreement (EULA) is shown the first time any user logs into the SL1 system, for users that use the new user interface as well as the PHP user interface.
- Updated the look and functionality for some drop-down lists to include a search field and multi-select checkboxes and to display selected fields at the top of the list.
- Updated the look and functionality for date and time selectors.
- Updated the look and functionality of the "share with" drop-down menus that are used to determine whether maps and dashboards are public, private, or shared with one or more specific organizations.
- Updated the look and functionality of the wizards used in SL1 to guide users through a process, such as guided discovery.

Issues Addressed in 10.2.0

Access Hooks and Access Keys

- Updated the DEV_THRESHOLDS_RETENTION Access Hook to ensure that it properly grants access to the Device Retention Thresholds setting, and that the setting is disabled for users who do not have that Access Hook aligned to their user profile. (Case: 00123026) (JIRA ID: EM-39704) (JIRA ID: EM-39644) EM-39704

Agent

- Addressed an issue with the SL1 Agent that affected network interface inventory and performance monitoring. (CASE: 00105708) (JIRA ID: EM-38798)

Alerts

- When using the active() function to define a "healthy" alert, if the alert_ID argument for the active() function specifies an alert with a very long index value, the healthy alert is now successfully triggered. (Case ID: 00098207) (Support ID: 110727) (JIRA ID: EM-26723)

Backups

- SL1 now supports full backups and configuration backups on Database Servers configured for Disaster Recovery. (Case: 00070215) (Case: 00055834) (JIRA ID: EM-36976) (JIRA ID: EM-38366) (JIRA ID: EM-39510) (JIRA ID: EM-33473)

- Added a new feature to backups. User can specify the number of backups to retain before new backups overwrite the older backups. (JIRA ID: EM-19761) (JIRA ID: EM-34904)
- Configuration backups now includes Business Service information. (Case: 00115987) (JIRA ID: EM-39122) EM-39122

Business Services

- Addressed an issue with Business Services topology map relationships. SL1 no longer displays an error message in the System Log when generating topology relationships in Business Services. (Case: 00108809) (JIRA ID: EM-38632) (JIRA ID: EM-38556)
- Addressed an issue in which Health and Availability values were not being calculated in business services with more than 10 device services. (Case: 0011805) (JIRA ID: EM-39289) (JIRA ID: EM-39230)

Collection

- To prevent the process "Data Collection: OS Process Check" from generating a System Failure message, added a timeout value to the internal SNMP walk function. If a single device is slow to respond to the "Data Collection: OS Process Check" process, SL1 will generate a warning message and continue to the next device. (Case: 00051486) (JIRA ID: EM-32811)
- Data Collectors no longer return the spurious SSL error "connecting to host using weakened cipher". SL1 now displays detailed log and error messages when using SSL connections. (Case: 00091230) (EM-36627) (EM-36547)
- In SL1 Extended Architecture where the Collector Pipeline is enabled, interface collection no longer causes multiple exceptions and no longer causes the Kubernetes pod "interface-store" to be restarted. (Case: 00114131) (JIRA ID: CPL-462)

Collector Pipeline

- In SL1 Extended Architecture where the Collector Pipeline is enabled, interface collection no longer causes multiple exceptions and no longer causes the Kubernetes pod "interface-store" to be restarted. (Case: 00114131) (JIRA ID: CPL-462)

Credentials

- Addressed an issue that was causing errors to appear when users attempted to scroll through the list of Organizations when creating a new credential. (EM-37769)
- Addressed an issue that was causing errors to appear when users attempted to scroll through the drop-down list of Collectors in credential tests and guided or unguided discovery workflows. (EM-37775)
- Addressed an issue in which Basic/Snippet credentials were not translating the variable %N as the device hostname, as intended. Now, when a user inserts the %N variable in the Hostname/IP field, SL1 will replace the variable with the hostname of the device using the credential. (EM-19364) (EM-19364)

Dashboards

- In the unified UI, in Dashboards, the Devices Widget no longer displays duplicate devices. (Case: 00099686) (JIRA ID: EM-37561) (JIRA ID: EM-37409)
- Addressed an issue in which the "ANY" filter was not working properly when searching for dashboards on the Dashboards page. (Case ID: 00111217) (Case ID: 00112066) (JIRA ID: EM-38780)

Database Server

- The em7_mailparse and postfix programs run on each Database Server, and postfix no longer causes the cluster to hang. (EM-37458)

Devices

- The filter-while-you-type fields now work correctly for the Device Hostname field on the Device Manager page (Devices > Device Manager or Registry > Devices > Device Manager in the classic user interface). (Support ID: 92976) (Case: 00090326) (Case: 00041082) (JIRA ID: EM-32028) (JIRA ID: EM-36468) (JIRA ID: EM-7926)

Dynamic Applications

- Addressed an issue where you could not create a presentation Object in a Dynamic Application that has the Show as Percent field set to NO. (Case 00045183. JIRA ID EM-32606)
- Addressed an issue in Dynamic Applications that was causing the "Align if OID is NOT Present" discovery object alignment condition field to not work as intended when the "Tabular" option was also selected. With this update, SL1 will ignore the "Tabular" option if the "Align if OID is NOT Present" field is also set. (Case: 00070409) (JIRA ID: EM-34944)

Email

- SL1 now supports all ASCII characters 0-127, including the dollar sign (\$) in email addresses. (Case: 00082451) (JIRA ID: EM-35966)

Events

- Addressed an issue that was causing fatal PHP error messages to appear when users attempted to clear an already cleared event. These error messages no longer appear in the PHP logs in that scenario. (Case: 00095972) (JIRA ID: EM-37340) (JIRA ID: EM-37124)
- Addressed an issue with Event Masks. When Event Mask is enabled with "Group in blocks" greater than 1 day, all events that occur during that time period on that device are now correctly grouped under a single event in the Event Console. (Case: 00069882) (JIRA ID: EM-35593)
- SL1 now triggers alert 290 "Process Time Exceeded" (Minor severity) and alert 291 "Process Time Exceeded" (Major severity) only once every minute, to avoid overloading the SL1 system. These alerts are used by two events: "System Process running longer than expected" with the event ID of 3336 and severity Minor and "System Process running longer than expected" with the severity Major and the event ID of 3337. (Case: 00048419) (Case: 00015712) (JIRA ID: EM-28131)

High Availability and Disaster Recovery

- The ScienceLogic Support PowerPack now generates a major event if the CRM configuration or corosync configuration are out-of-date. (Case: 00022805) (Case: 00026718) (JIRA ID: EM-30285) (JIRA ID: EM-40476)

IP Address

- After changing the IP address of an SL1 appliance in the Web Configuration Utility, both the Web Configuration Utility and the System Status script both accurately display the SL1 appliance as "licensed". (Case: 00100819) (JIRA ID: EM-37445) (JIRA ID: EM-37444) (JIRA ID: EM-37318)

Maps

- Addressed an issue with LLDP topology maps. While discovering and generating LLDP topology maps, SL1 no longer generates an unhandled exception if the SNMP OID value for the chassis is "None". (Case: 00061034) (Case: 00057421) (Case: 00007070) (JIRA ID: EM-28243) (Support ID: 172487)

MariaDB

- This release includes an upgrade to MariaDB. The upgrade includes jemalloc as the memory allocation tool for MariaDB. 10.1.x included an upgrade to MariaDB that did not include jemalloc, and the upgrade caused out-of-memory errors. SL1 10.1.5.3 and SL1 10.2.0 add jemalloc back to the SL1 system.
 - This issue affected:
 - 10.1.0
 - 10.1.1
 - 10.1.2
 - 10.1.3
 - 10.1.4
 - 10.1.4.1
 - 10.1.4.2
 - 10.1.5
 - 10.1.5.1
 - For SL1 versions 10.2.0 and later, jemalloc is included with the platform.
 - For SL1 versions prior to 10.1.0, jemalloc is included with the platform.

MIB Compiler

- Addressed a typo in the heading of the MIB Compiler page (System > Tools > MIB Compiler) in the 8.4.2 version of the platform. (Support ID: 139177, Case: 00099142) (JIRA ID: EM-21328)

PhoneHome Collector

- The administration password for a Data Collector is no longer reset to default after establishing a connection and syncing between the Database Server and a Data Collector in a PhoneHome configuration. (Case: 00104590) (JIRA ID: EM-37987)
- To prevent the error "Operation not permitted: '/var/log/phonehome/shell_phonehome.log'", the PhoneHome shell logs are now stored in /home/\$USER/logs/shell.log of each each PhoneHome Database Server and each PhoneHome Data Collector. SL1 auto-rotates these log files, to prevent the log file from filling the /home partition. (JIRA ID: 40007) (Case: 00148549) (Case: 00145094) (Case: 00135470)
- SL1 now successfully auto-rotates the log files in /var/log/phonehome. JIRA ID: EM-39054) (Case: 00115632) (Case: 00132579) (Case: 00115632)
- Phonehome SSH tunnels (em7_ph_tunnels service) no longer shut down during MariaDB upgrade or during temporary loss of connection to MariaDB. (JIRA ID: EM-40512) (Case: 00134822)
- The command "phonehome clear all" (if run from the primary Database Server) now displays a detailed description and warning before deleting the PhoneHome configuration on the Primary Database Server, secondary Database Server, and DR Database Server and the tunnels for each Data Collector, Message Collector. The "phonehome clear all" (if run from the secondary Database Server) now displays a message discouraging users from using the command. (JIRA ID: EM-40230) (JIRA ID: EM-41449)
- Improved the PhoneHome check command and the PhoneHome msg command to generate cleaner stdout and prevent the error "Error: Problems when converting string to json - No JSON object could be decoded". (JIRA ID: EM-39994) (JIRA ID: EM-40020) (Case: 00126012)
- PhoneHome will not try to change the shell for the phoneHome users in order to force a reconnect from the client. Rather, it will not resort to killing client processes to force a reconnect. (Case ID: 00125995) (Case ID: 00134547) (JIRA ID: EM-41028)
- PhoneHome does not escalate privileges for the watchdog functionality on the collectors causing pam_unix critical errors. (Case ID: 00057979) (Case ID: 00061414) (Case ID: 00104655) (Case ID: 00113365) (JIRA ID: EM-37215)
- PhoneHome collectors no longer need to elevate privilege using sudo to check for tunnel connectivity. (Case: 00150558) (Case: 00137119) (Case: 00134802) (Case: 00125446) (Case: 00087715) (JIRA ID: EM-35596)

Platform

- The SL1 system process "EM7 Core: Task Manager" (proc_mgr.py) no longer blocks all processes when a single process fails. (Case: 00070402) (Case: 00068102) (Case: 00067779) (Case: 00074213) (Case: 00074804) (Case: 00080930) (Case: 00042477) (JIRA ID: EM-31636) (JIRA ID: EM-35226)
- The SL1 process "EM7 Core: Daily Maintenance" (maint_daily.py) no longer triggers an unhandled exception and then triggers a long-running database transaction. (Case: 00102461) (Case: 00091288), (Case: 00091639), (Case: 00091644), (Case: 00093043), (Case: 00093491), (Case: 00093687) (JIRA ID: 37006)
- For the SL1 Process "EM7 Core: Daily Maintenance" (maint_daily.py) refactored the process that prunes data for Configuration Dynamic Applications to decrease memory usage. (Case 00010389) (Case: 00014808) (Case: 00036875) (Case: 00071766) (Case: 00073612) (Case: 00092767) (Case:

00092780) (Case: 00093347) (Case: 00093365) (Case: 00094290) (Case: 00094292) (Case: 00096260) (JIRA ID: EM-20781)

- To prevent the process "Data Collection: OS Process Check" from generating a System Failure message, added a timeout value to the internal SNMP walk function. If a single device is slow to respond to the "Data Collection: OS Process Check" process, SL1 will generate a warning message and continue to the next device. (Case: 00051486) (JIRA ID: EM-32811)
- The process "Data Collection: Dynamic Refresh" (dynamic_check.py) no longer causes DEADLOCK failures for storage objects. (Case: 00092599) (JIRA ID: EM-36916) (JIRA ID: EM-39683)
- Addressed an issue with SQL queries. Queries with more than 1000 entries in the IN condition no longer return incorrect results. (Case: 00131447) (JIRA ID: EM-39632)
- SL1 now includes automated self-healing. Self-healing jobs can (Case: 00097130) (JIRA ID: EM-37874) (JIRA ID: EM-37454):
 - automatically set "s-em7-core:s-em7-core" as the owner and group for the file "silo.log"
 - automatically restart the process em7_patch_manager if it is stuck in deactivating mode
 - automatically kill queries that have run for longer than 1 hour and logs each query

PowerPacks

- Addressed an issue where disabling the em7admin user (UID: 1) prevented powerpack_batch_install.php from executing, and the Update column did not appear on the PowerPack Manager page (System > Manage > PowerPacks during an SL1 upgrade. (Case: 00102311) (JIRA ID: EM-37628)

Process Monitoring

- Addressed an issue in Process Monitoring that was causing the process instance count to not work as expected when some processes had arguments and some did not. With this update, the process arg_regex pattern now counts processes without arguments when the arg_regex is provided as '.*'. (Case: 00032340) (JIRA ID: EM-31367)

Reports

- Addressed an issue with the Event Detections Report where the report contained incorrect times in the "Last Detected" or "First Occurrence" fields. Version 1.6 of the report is included in the latest Core Reports PowerPack. (Case: 00099302) (JIRA ID: EM-15574)
- Addressed an issue with the Event Detections Report where the report did not show all active events for the devices specified in the report. (Case: 00030706) (JIRA ID: EM-30346)
- Addressed an issue with the Event Detections Report where the report contained incorrect event counts. (Case: 00090325) (JIRA ID: EM-36467)
- Addressed an issue with the Asset List Report where the report showed all assets instead of filtering the report by Service Status. (Case: 00011274) (JIRA ID: EM-27631)
- Addressed issues with Devices > Performance Multi Device Reports where the reports did not show all devices specified for the report. (Case: 00107129) (JIRA ID: EM-38759)

- Removed the HTML image files linked to HTML reports in the /opt/em7/gui/ap/www/em7/libs/od_templates/populated directory from daily maintenance tasks. This update prevents this directory from being cleaned in daily maintenance tasks, as SL1 now maintains report files with the new retention policy. (Case: 00061199) (JIRA ID: EM-33720)
- Addressed an issue where a scheduled report was terminated without completing (sigterms), and the temporary files were left in the /tmp directory, causing that directory to run out of space. (Case: 00088452) (JIRA ID: EM-24360)
- Addressed an issue that occurred when an Administrator user deleted a user account, but all scheduled reports belonging to that user were not be removed. (Case: 00098638) (JIRA ID: EM-22527)

REST API

- Addressed an issue where the API call for ticket_log timed out. (Case 00067563. JIRA ID EM-34403)

Security

- A library was updated to address a potential cross-site scripting vulnerability. (Case: 00101374) (JIRA ID: EM-38886)

Ticketing

- Addressed an issue where a user could not delete a ticket that had a Status of "Resolved". (Case 00065433. JIRA ID: EM-20713)
- Addressed an issue with Ticketing that was preventing external tickets from displaying when users clicked on incidents in the Device Summary tab. (Case: 00016778) (Support ID: 174828) (JIRA ID: EM-28488)
- Addressed an issue where double-spacing and tabbing caused special characters to display in a ticket, usually when using a Chrome browser. (Case 00054531, 00057820, 00048275, 00003494. JIRA ID EM-28892)
- Addressed an issue where SL1 added extra characters to a ticket after the ticket was saved. (Case: 00039765. JIRA ID: EM-31261)

Topology Maps

- Addressed an issue with LLDP topology maps. While discovering and generating LLDP topology maps, SL1 no longer generates an unhandled exception if the SNMP OID value for the chassis is "None". (Case: 00061034) (Case: 00057421) (Case: 00007070) (JIRA ID: EM-28243) (Support ID: 172487)

User Policies

- Addressed an issue where a user could not disable the Time Zone field in the User Policy Properties Editor (Registry > Accounts > User Policies > select a policy). (Case 00114900. JIRA ID EM-38104)

Windows

- Added missing dependencies for Python modules that support winrm. (Support ID: 174528) (Case: 00099831) (Case: 00014412) (JIRA ID: EM-28227)

User Interface

In 10.2.0, the default user interface is the Classic User Interface. To change the default user interface to the Unified UI:

To change the default user interface to the Unified UI:

1. Go to the console of the Administration Portal or open an SSH session to the Administration Portal.
2. Navigate to `/opt/em7/share/config/nginx.d`
3. At the shell prompt, enter the following:

```
sudo vi em7ngx_web_ui.fragment
```

4. Find this section in the file:

```
location / {
    root /usr/local/silo/gui/ap/www;
    index index.em7;
}
```

5. Edit as follows:

```
location = / {
    proxy_set_header Host $host;
    proxy_set_header X-Forwarded-For $remote_addr;
    proxy_set_header X-Forwarded-Proto $scheme;
    proxy_connect_timeout 10;
    proxy_read_timeout 300;
    proxy_pass http://localhost:3000;
}

location /em7 {
    server_name_in_redirect off;
    root /usr/local/silo/gui/ap/www;
    index index.em7;
}
```

6. Save your changes to the file (`:wq`).
7. Restart the web server. To do this, enter the following at the shell prompt:

```
sudo systemctl restart nginx
```

SL1 Extended Architecture

10.2.0 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 10.2.0

Before upgrading to 10.2.0, 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.

The 10.2.0 release includes the following PowerPacks that are new or updated and included with the release:

- Cisco: UC Ancillary v103
- Cisco: UCS Standalone Rack Server v103
- Microsoft Hyper-V v101
- Microsoft: Azure v112

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 ISO to comply with the new SL1 pricing model:

NOTE: If you are upgrading from a previous version of SL1, the 10.2.0 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 10.1.4.

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:

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](#).

1. Log in to the [ScienceLogic Support](#) site. 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 are interested in and click its name:

Release Version
SL1 Colosseum 10.1

Edit
Printable View
Delete
▼

Release Name

SL1 Colosseum 10.1

Version

10.1

End of Maintenance

12/31/2021

Latest Release Date

7/2/2020

Online Documentation

<https://docs.sciencelogic.com/10-1-0/>

End of Life

6/30/2022

Allow Customers to View on Community

☒

Release Files (6+)

New

File Name	Comments	Record Type	Release Date
10.1.5.3	8.12.1.3 introduced del...	Product Update	3/29/2021
10.1.5.2	[LA Release] 8.12.1.3 in...	Product Hotfix	3/18/2021
10.1.5.1	8.12.1.3 introduced del...	Product Update	2/4/2021
10.1.5	8.12.1.3 introduced del...	Product Update	1/26/2021
10.1.2.2	[LA Release] 8.12.1.3 in...	Product Hotfix	12/28/2020
10.1.3.2	[LA Release] 8.12.1.3 in...	Product Hotfix	12/28/2020

View All

Release Announcements

8.14.6 VERSION UPGRADE FAILURE / 21-08-20

In 8.14.6, ScienceLogic introduced an issue that causes System Update to fail on SL1 systems that do not use the UCAPL upgrade. For additional details, please read [this article](#). This issue will be addressed in the upcoming 8.14.7 release.

8.12.0 OR 8.12.0.1 / 05-10-19

In 8.12.0, ScienceLogic introduced an issue that caused System Update to fail on any SL1 appliance that contains 19 or more CPU cores. This issue will be addressed in the upcoming 8.12.0.2 release.

If your SL1 system includes one or more appliances with 19 or more CPU cores, please read [this article](#) and perform the steps in the article during installation of 8.12.0

Release Announcements Archive

- 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.
- Under **Files**, select the link for the file you want to download. The file is then downloaded to your local computer.

Recently Deprecated Features

8.9.0

- High-Availability and Disaster Recovery no longer support the command "drbd-overview". To check the status of drbd, use "cat /proc/drbd".

8.9.2

- Added security features to the ScienceLogic API. The /api/account resource no longer includes the "passwd" field. To set a password for a user account, administrators can POST to "/api/account/<account_id>/password". The password value will be encrypted in storage. (EM-15299)

- The Knowledge Base includes known security vulnerabilities. ScienceLogic no longer supports the Knowledge Base and strongly recommends that users disable the Knowledge Base. In future versions of SL1 will, the Knowledge Base will be disabled by default. (EM-26508)

8.10.0

- Deprecated the Access Hook "Cred:Passwords". Added two new Access Hooks to provide more granular access: "Cred: Passwords: Edit" and "Cred: Passwords: View". (EM-27312)

8.12.0

- Removed the password field from the account resource in the ScienceLogic API (EM-26716)
- The FTP, SFTP, NFS, and SMB backup options that stage locally are no longer supported. (EM-28362)
- Integration Server appliances are no longer supported. (EM-27126)
- System Update no longer supports the shell command "deploy_patch". (EM-23982)

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)

Known Issues

- If you are using SL1 Extended Architecture and have enabled Machine Learning for one or more devices, there is a Known Issue. If you have changed the default administrator's username or password, there is a Known Issue when starting the services for Machine Learning. The remediation steps are in this Knowledge Base article:
https://sciencelogic.lightning.force.com/lightning/r/Knowledge_kav/ka04z000000HCczAAG/view
- ScienceLogic strongly recommends that you review all [Known Issues](#) for SL1 (<https://support.sciencelogic.com/s/topic/0TO0z000000E6w7GAC/known-issues>).

© 2003 - 2021, 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



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

International: +1-703-354-1010