



---

# SL1 10.1.0 Release Notes

SL1 version 10.1.0, Rev 6

---

# Table of Contents

<b>Overview</b>	<b>1</b>
Upgrade Notes	1
Caveats	4
User Interface	5
SL1 Extended Architecture	5
New Features in 10.1.0	6
Issues Addressed in 10.1.0	12
PowerPacks in 10.1.0	16
New and Updated Packages in 10.1.0	17
Beta Features	18
Disabling the Knowledge Base	18
Upgrade Process for Systems Running 8.1.0 and Earlier	19
Upgrade Process for Systems Running 8.1.1 and Later	19
Downloading SL1 Updates on SL1 Systems running 8.1.x - 8.5.x	20
Downloading SL1 Updates on SL1 Systems Running 8.6.0 or Later	20
Upgrade Process for SL1 Extended Architecture	21
Special Steps for Upgrading MariaDB	21
Manually Updating ap2	23
Recently Deprecated Features	24
8.9.0	24
8.9.2	24
8.10.0	25
8.12.0	25
8.14.0	25
10.1.0	25
Future Deprecations	25
Known Issues	26

---

# Overview

This document describes:

- [Upgrade Notes](#)
- [Caveats](#)
- [User Interface](#)
- [SL1 Extended Architecture](#)
- The [New Features](#) in the 10.1.0 release
- The [Issues Addressed](#) in the 10.1.0 release
- The [PowerPacks](#) included in the 10.1.0 release
- [New and Updated Packages](#) in the 10.1.0 release
- [The Beta Features available with the 10.1.0 release](#)
- [Disabling the Knowledge Base](#)
- [Upgrade Process for Systems Running 8.1.0 and Earlier](#)
- [Upgrade Process for Systems Running 8.1.1 and Later](#)
- [Upgrade Process for SL1 Extended Architecture](#)
- [Special Steps for Upgrading MariaDB](#)
- [Manually Updating ap2](#)
- The [recently deprecated features](#)
- The [features deprecated](#) in the next release
- The [Known Issues](#)

**TIP:** 10.1.0 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 10.1.0.

---

## Upgrade Notes

**CAUTION:** The 10.1.x line requires additional steps when upgrading MariaDB. For details, see the section [Special Steps for Upgrading MariaDB](#), or see the chapter on [Upgrading SL1](#) in the [System Administration](#) manual or view that chapter [online](#).

- **If you are running a version prior to 8.12.0, do not install 10.1.0 if you do not plan to immediately consume 10.1.0.** After you import the 10.1.0 release, all appliances in your SL1 system will now use the

**new system update.** After you import the 10.1.0 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.

- 10.1.0 includes important security updates. **You must reboot all AIO appliances and all appliances in the SL1 Distributed Architecture after deploying 10.1.0.** If you would like assistance in planning an upgrade path that meets your security needs while minimizing downtime, please contact your Customer Success Manager.
- 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 this section online: [https://docs.sciencelogic.com/latest/index.htm#Web\\_Admin\\_and\\_Accounts/System\\_Administration/sys\\_admin\\_system\\_upgrade.htm](https://docs.sciencelogic.com/latest/index.htm#Web_Admin_and_Accounts/System_Administration/sys_admin_system_upgrade.htm))
- 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.
  - 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.
  - 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 8.10.3 to SL1 8.14.3, you would install only a single update. And looking at the table below, you can see that 8.14.3 requires you to reboot all SL1 appliances after upgrade but does not require you to upgrade MariaDB. However, two of the releases between 8.12.0 and 8.14.0 include an upgrade to MariaDB. Therefore, you must reboot all SL1 appliances and upgrade MariaDB after you upgrade to 8.14.3.

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

SL1 Release	Requires Reboot?	Requires MariaDB Upgrade?
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	No
8.14.9	Yes	No
10.1.0	Yes	Yes

## Caveats

Consider the following caveats before deploying 10.1.0:

- The 8.14.6 release includes changes that are available in the upcoming 10.1.1 release but are not available in the 10.1.0 release.
- Changes from the 8.14.5 release are supported in 10.1.0. You can upgrade from 8.14.5 to 10.1.0.
- 10.1.0 does not support Data Collectors running earlier versions of SL1 if the SL1 Database Server is 10.1.0. You must upgrade all SL1 appliances to 10.1.0, using the System Update tool (System > Tools > Updates).
- As of 8.12.2, ScienceLogic will no longer update 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 spawn a new window that displays the appropriate page from [docs.sciencelogic.com](https://docs.sciencelogic.com).
- **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 10.1.0.***
- 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](#).***

- 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-upgrade tasks. The post-patch update process is automatically stopped after 24 hours. However, depending on the size of your database as well as the version from which you are upgrading, the post-upgrade tasks can take several days to perform. If the post-patch update process is stopped after 24 hours, the process will automatically re-start and continue processing from the point at which it was stopped. If you see an event that indicates the post-patch update process was stopped, you do not need to contact ScienceLogic support for assistance until you see the same event for three consecutive days.

---

## User Interface

In 10.1.0, the default user interface is the Classic User Interface.

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

10.1.0 includes the SL1 Extended Architecture.

SL1 includes one or more appliances or nodes (part of a cluster) that function together to provide the SL1 platform and application. There are three primary configurations:

- **All-In-One**. In this configuration, a single appliance provides all the functions of SL1. The capacity of an All-In-One instance cannot be increased by adding additional appliances. This configuration is best for smaller deployments.
- **Distributed**. In this configuration, the functions of SL1 are divided between multiple appliances. A **Distributed** instance of SL1 can be as small as two appliances or include multiple instances of each appliance. This configuration is best for production environments that monitor lots of devices or that monitor a large volume of data for each device.
- **Extended**. An extension of a Distributed instance. The **Extended** configuration adds both a **Compute Cluster** and a **Storage Cluster**. The Compute Cluster includes multiple Compute Nodes. The Storage Cluster includes multiple Storage Nodes. The **Extended** configuration also adds a Management Node to install and update the Compute Cluster and Storage Cluster, and one or more Load Balancers to manage the workload to the Compute Cluster. This configuration provides scale and can take advantage of the SL1 Agent to collect detailed data about devices and applications.

Resiliency and redundancy can also be accomplished by adding additional appliances to these configurations.

10.1.0 allows users to upgrade an existing Distributed system, upgrade an existing Extended system, and add the SL1 Extended Architecture to existing systems.

For upgrade details, see the section on Upgrading to 10.1.0

For instructions on how to deploy the SL1 Extended Architecture, see the *Installation* manual.

---

## New Features in 10.1.0

To ensure that you are using the "unified" SL1 user interface, type **/ap2** at the end of the URL or IP address for your SL1 system. For example, you could type **<https://sl1.sciencelogic.com/ap2>**.

### Activity Center

- A new Activity Center has been added that enables you to run a set of generic network diagnostic tools for devices and also enables you to run a set of generic network diagnostic tools and custom user-initiated actions for events. You can also view a history of tools that have been run for a device. You can access the Activity Center in one of the following ways
  - Clicking the Activity button in the navigation bar at the top of any page in SL1
  - Clicking the Tools button on the Device Investigator page for a particular device
  - Expanding the Device Drawer for a device and running a tool from the Tools pane
  - Expanding the Event Drawer for an event and running a tool or action from the Tools pane
  - Opening the Event Investigator page for an event and running a tool or action from the Tools pane

- Opening the Service Investigator page for a business service, expanding the Event Drawer for an event on the Events tab, and running a tool or action from the Tools pane

## Agent

- The installation pages for the SL1 Windows Agent and SL1 Linux Agent were updated.
- The SL1 Windows Agent v128 supports UTF-8 characters (specifically, Japanese).
- You can now install the SL1 Agent on a Solaris server. This is a beta feature in 10.1.0.
- You can now install the SL1 Agent on an AIX server. This is a beta feature in 10.1.0.
- For each device monitored with the SL1 Agent, you can now manage the SL1 Agent configuration settings on the Settings page of the Device Investigator.
- The Agents page (Devices > Agents) was updated to include Agent Version, Newest Version, and Upgrade options.
- You can now perform bulk actions on agents on the Agents page (Devices > Agents).

## Business Services

- The Overview page in the Business Service Investigator has been redesigned to incorporate behavioral service event correlation and anomaly detection. The redesigned page displays key metrics, events, and anomalies that are impacting the selected service, and will enable you to view, understand, and troubleshoot service problems from a single page.

## Collector Pipeline

- SL1 now includes a new architecture feature, Collector Pipeline. This pipeline works with the extended architectures and allows you to horizontally scale your SL1 configuration and improve the performance of data collection.

## Credentials

- The pages used for managing and creating/editing credentials have been redesigned in SL1.
- You can now test credentials during the guided or unguided discovery processes or from the Credentials page (Manage > Credentials). In addition, the credential test modal pages that appear when you use these methods have been redesigned.
- What does a networking seal say?
- Arp! Arp! Arp!

## Dashboards

- In the classic user interface, the (base) Context Topology Map widget now uses HTML5 instead Flash.
- In the classic user interface, the "Topology Map" (non-base) uses HTML5 instead of Flash.
- When creating or editing dashboard widgets, you can now pick specific devices or services to always appear in the widget, in addition to any devices or services that might appear based on context subscription/driving or data filtering.



- In dashboards, users can now view data from all descendants of any service context in Device Leaderboard widgets. Additionally, in dashboard widgets that subscribe to a service context, users can now filter the list of subscriber services that appear in the widget. When creating sunburst widgets in service dashboards, users can now use a "Show devices" toggle field to determine whether or not devices appear in the outermost sunburst ring.
- For leaderboard widgets, you can use the "Select Index" metric option to choose which index to display from a particular metric.
- A new NOC Overview dashboard has been added to SL1. This dashboard is included in the SL1 system by default and displays a high-level overview of your business services and their current health statuses in a single-pane view. For more information about this dashboard, see the Dashboards manual.

## Devices

- In the classic user interface, the Topology tab in the Device Performance panel now uses HTML5 instead of Flash.
- You can now unmerge a physical device from its virtual component device when the merged device is in a vanished state.
- The Device Investigator page has been updated to automatically save any changes you make to the set of metrics that display for each device upon exiting the page, and to recall these changes when you next visit the page. You also have the option to apply those changes to every device in the device's device class, or to reset the changes to their default settings.
- Users now have an option within the Device Investigator to view a map that displays the relationships from a particular device.
- In the Device Investigator page, users can now change the order in which widgets appear on the page by dragging and dropping the panel headings on the left sidebar into their desired order. Users can also use this drag-and-drop method to combine timeseries metrics on the left sidebar so those metrics appear in combined charts on the page. This enables you to compare related metrics in the same chart.
- You can now add devices using guided or unguided discovery from the Discovery Sessions page.

## Events

- The Event Engine now includes a built-in timeout for Event Policies with a source "Email". If SL1 hangs due to a poorly formed RegEx in the event email, SL1 stop evaluating the event when the timeout value (10 seconds) is exceeded. When a timeout occurs, a System event will be generated. This prevents poorly formed RegEx from crashing the Event Engine.

## HA

- Improvement to High Availability configuration. The query for HA status no longer tries to access the RPM database to determine if the parent database is remote. The query now reads from a local file. If the local file does not provide the database status, the HA status query will try to access the RPM database.

## JMX

- The SL1 Extended Architecture allows users to align JMX Dynamic Applications to devices monitored by the SL1 Linux Agent.

## Kafka

- SL1 Extended now uses the Kafka Messaging service.

## Maps

- In the Classic UI, you can now choose between HTML5 or Flash as the default setting for maps. Use the "Map Rendering" field on the Behavior Settings page (System > Settings > Behavior) to define the default.
- The Expansion Depth field lets you determine the number of node tiers that are expanded by default when you create or edit a map, or add a node to a map
- You can add images, rectangles, circles, and other shapes and text to a map. These features let you illustrate how your devices are working within the context of their geographic location and how they are used.
- You can choose to display a color-coded legend on the map that shows node status or edge type.
- The following user interface elements were updated and enhanced:
  - interactive legend on the left side of the Maps window,
  - buttons for Zoom In, Zoom Out, Fit All, and Center,
  - Properties pane for a map.
- If you select two device nodes on a map, you can create a relationship between those two device nodes in the map. The relationship will only exist in this map, not anywhere else in SL1.
- New map layout option. The Multi-Parent Hierarchical layout best displays nodes that have a large number of parent-child relationships, which get arranged by edge direction. This layout is similar to the Classic Maps (Views) from the classic user interface. New Target button on the Properties pane for a map.
- Select a device or service on a map and then click the Locate icon to create a targeted map for the selected node. A new Map window opens, and the map has a default Expansion Depth of 1. This feature is useful if you want to focus on a small portion of a very large or detailed map.
- On the Classic Maps page, you can view HTML5 versions of maps that are based on the Classic Maps (Views). You can also continue to use Classic Views with Flash enabled. Use the "Map Rendering" field on the Behavior Settings page (System > Settings > Behavior) to choose between HTML5 or Flash as the default setting for maps.

## MariaDB

- 10.1.0 includes an upgrade from MariaDB 10.1 to MariaDB 10.4.

## PhoneHome

- Added a new utility to PhoneHome, called monitor\_phonehome, to help debug and diagnose issues with PhoneHome.

## Platform

- The database password in /etc/silo.conf is now encrypted.
- Added a new feature to logging. SL1 now includes logs that allow analysis of Data Pull Storage Object Processing.

- SL1 now includes a new architecture feature, Collector Pipeline. This pipeline works with the extended architectures and allows you to horizontally scale your SL1 configuration and improve the performance of data collection.
- SL1 Extended now uses the Kafka Messaging service.

## PowerShell

- Addressed an issue with PowerShell Dynamic Applications. Added a lock to mysql to prevent multiple collection processes from writing or overwriting changes to the /etc/krb5.conf file. (Support ID: 119845) (JIRA ID: POSH-36) (JIRA ID: EM-17823) (JIRA ID: POSH-127) (EM-18388)
- For PowerShell Dynamic Applications, now ensure that each exception in the log files includes the device ID and the application ID.

## Product Telemetry

- To eliminate duplicate devices during billing, the payload sent from an SL1 stack to the Billing Server includes a list of all relationships of all types. For each relationship, the following information is included: 1) The type of relationship. For DCM and DCM+R relationships, this is the GUID of the collection object ID used to identify the relationship; 2) The device class GUID for the device at the other end of the relationship.
- To improve accuracy, the payload sent to the Billing Server includes the count of appliances of each type monitored by your SL1 system.
- SL1 now includes a database settings that allows subscription billing to include vanished devices in the payload sent to the billing server.

## Publisher

- SL1 includes a new beta feature, Publisher. Publisher allows users to consume data collected by SL1 (via Data Collectors and the ScienceLogic Agent) in real-time, outside the user interface. Publisher can provide data for long-term storage or provide input to other applications the perform analysis or reporting.

## Reports

- Video reports were updated to include Cisco CE Series data.
- The following reports were added to the Devices category (Reports > Run Report > Devices): Blackberry Top-N Performance Dynamic App and Blackberry Top-N Report Vitals. The following report was added to the Network Interfaces category (Reports > Run Report > Network Interfaces): Blackberry Top-N Interface Statistics.
- Updated a number of Video Reports to address an issue with non-UTC time zones generating data for two days instead of one.
- When you create or edit a report, you can use SQL queries to help SL1 estimate the number of rows in the report before generating the report. You add one or more SQL queries to the Query field on the Report Template Editor page. The result of the additional query or queries appears in the Row Count Estimate field for the report. You can add up to eight queries.
- As you update the fields for a custom report, the Row Count Estimate field next to the Generate button is updated to provide an estimate of the number of rows that will appear in this report. You can use this field to manage the size of the generated report by adding or removing items from the report as needed. For more

information about how to edit the settings for this field, see "Creating a Report Template" in the Report Development manual

- The Row Count Estimate field updates each time the user selects or de-selects Inputs, such as selecting Devices, then de-selecting Devices and selecting Assets instead. To disable the Row Count Estimate field, de-select the Report Size Estimation option on the Behavior Settings page (System > Settings > Behavior).
- The following reports utilize the Row Count Estimate field: Device-At-A-Glance, Device Availability, Device Top Metrics, Device Top Utilization, Device Utilization, Interface In Use, Interface Top Metrics, Monitored Element, and PowerPack Information.
- The EM7 Core Reports was updated to version 111 to include Report Size Estimate queries for reports.
- Quick Reports run on the current appliance (the appliance where the user is logged in) in a distributed SL1 stack.
- When creating a report job, you can use the Report Appliance field to select the appliance on which to run the report. By default the current appliance is selected. If you select "Active Database", SL1 will locate the active database server at run time and will execute the report on that database server. If you are running a Quick (ad hoc) Report, SL1 executes the report on the current system you are logged into.

### **Run Book Automation**

- Run Book Automations includes a new feature that allows user-initiated Run Book Actions. You can manually trigger automations via an Event for a Device. You can also define Run Book Automation policies that can be triggered by an event and can also be triggered manually by a user. New PowerPacks with the words "user-initiated" or "automation" in the title include these new Run Book Automations.
- SL1 now includes a new access hook, EVT\_AUTOMATION\_USER\_INITIATED, that allows users to execute user-initiated Run Book Automations. The new access hook is aligned to "Grant All" key by default.
- SL1 includes a GQL interface for user-initiated Run Book Automations.

### **Search**

- You can now save searches for business services and maps. In addition, saved searches now appear at the bottom of the drop-down menu that appears when you begin typing in the Search field.

### **SMTP**

- Added a backend setting that allows you to tune the timeout value for the SMTP Server connection.

### **SNMP Collection**

- Added a new beta 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.
- Users can collect metrics for new beta feature, Concurrent SNMP Collection.

### **SSH**

- SL1 now includes a new SSH collection process. Several upcoming PowerPacks will take advantage of this service to improve scale.

### **System Update**

- Multiple improvements to System Update including a suite of pre-upgrade checks with feedback to the user that determines if your SL1 stack can be successfully upgraded.

### Themes

- Users can now create and deploy user interface themes in the new SL1 user interface. Users with the proper permissions can also edit themes, add themes to a PowerPack, assign hostnames to themes, and create custom login pages for default system themes. Themes set in the new SL1 user interface will also carry over to the unified SL1 user interface.

### User Interface

- Users are no longer required to have the AP\_ACCESS hook to view the new SL1 user interface.

### Views

- In the classic user interface, the Views tab uses HTML5 instead of Flash.

---

## Issues Addressed in 10.1.0

### Agent

- The UI now includes a working IP address for downloading the Agent. (Case: 00067793)

### Access Logs

- Users that access SL1 through SAML SSO are now captured properly in the Access Logs (System > Monitor > Access Logs). (Case: 00017201) (JIRA ID: EM-32037)

### Authentication

- The SecurID passcode presented to a user when the Authentication Profile uses MFA no longer limits the input to 6 characters. (Case: Case 00082425) (JIRA ID: EM-35555)
- Fixed a security vulnerability that allowed a user to manipulate the URL to switch between logins in the classic user interface and the next user interface after providing valid credentials but before providing the MFA passcode. (Case: Case 00082425) (JIRA ID: EM-35617)

### Dashboards

- Addressed an issue with index labels that was causing no data to appear in dashboard Top-N widgets, even when that data was present and being collected in the monitored system.
- In Dashboard widgets that cannot be successfully displayed in Detach mode, the Detach entry no longer appears in the Options menu. (Case: 00081900) (JIRA ID: EM-35236)

### Device Maintenance

- Updated the behavior for enabling and disabling scheduled device maintenance. Now, when you disable an active maintenance schedule, the device will be pulled out of maintenance mode; if you re-enable the schedule during the scheduled time window, the device will return to maintenance mode.

## Devices

- An issue was addressed that was causing the Merge Devices page to not load.

## Events

- Addressed an issue with the Dynamic Application "EM7: Event Statistics". After fixing a pathname, the Dynamic Application now returns values. (Case: 23530), (Support ID: 107152), (Support ID: 152756), (Support ID: 152756), (Support ID: 90047) (JIRA ID: EM-29557) (JIRA ID: EM-23482)
- When generating an event for a device group via snippet: generate\_alert(), the resulting event now includes a value for xname and the name column in the Event Console displays a value. (Case: 00006175) (JIRA ID: EM-30911)
- Events with a source of "syslog" or a source of "trap" now recognize the period (.) character as a newline character. (Case: 00064358) (JIRA ID: EM-34333)
- SL1 no longer throws an unhandled exception and stops processing Events of type API when an Event of type API is triggered for an IPv4 Network. (Case: 00066705) (JIRA ID: EM-34720)
- If an Event Policy has an "Occurrence Count" of 1 and "Occurrence Time" set to the Dynamic Application poll time, the event will trigger on the first occurrence of the event conditions, instead of the second occurrence. (JIRA ID: EM-35852) (Case: 00084756)

## HA/DR

- Addressed an issue with the HA/DR documentation. (Case: 00048761) (JIRA ID: EM-32738)
- Addressed an issue with High Availability configurations. SL1 now successfully stops Postfix, allowing the High Availability databases to fail over when necessary. (Case: 00061918) (JIRA ID: EM-33862)

## IT Services

- An issue was addressed that was causing IT Service maintenance schedules to end prematurely when more than one maintenance schedule overlapped.
- An issue was addressed that was causing the query functionality to not work as designed when creating an IT service.

## Licensing

- SL1 no longer generates a spurious event that says that the Administration Portal is not licensed. (Case: 00101336), (Case: 00086522), (Case: 00085097), (Case: 00070398), (JIRA ID: EM-34606)
- The licensed state for Administration portals is now correctly gathered by the Support Power-Pack. (Case: 0008157) (Case: 00103055), (Case: 00102649), (Case: 00092577), (Case: 00092666), (Case: 00088091), (Case: 00084387), (Case: 00081572) (JIRA ID: EM-35267)

## Network Interfaces

- All interface-related alert messages now contain the ifIndex value.
- Increased the DNS caching TTL from 60 seconds to 300 seconds, to prevent the process "Data Collection: Interface Bandwidth" from throwing unhandled exceptions.

- Addressed an issue with data collection for interfaces that caused the unhandled exception `"/opt/em7/backend/silo_collect/collect_if.py", line 1134, in do_validation...`.

## Performance Graphs

- Addressed an issue that was causing interface percentile graphs to erroneously show values higher than the interface speed.
- Addressed an issue that was causing some performance data values to appear as octets when they should have appeared as bits per second.

## PhoneHome

- If you are using PhoneHome Message Collectors in a distributed architecture, you can configure system settings to prevent the Data Collectors from incorrectly setting the URL for Gen 1 agent installations as the loopback IP address (instead of the local IP address for the Data Collector). (Support ID: 137962) (Case: 00007816) (Case: 00098940) (Case: 00055631) (JIRA ID: EM-21416) Case: 00098385) (Support ID: 161982) (JIRA ID: EM-24985)

## PHP Developer Logs

- Added a new feature to PHP Developer Logs. You can now download logs from the PHP Developer Logs page (System > Tools > PHP Developer Logs) in the classic user interface.

## Platform

- The script `/opt/em7/share/scripts/generate-silo-conf.py` no longer changes the IP address specified in `/etc/silo.conf` to the loopback IP. (Case: 00098179) (Support ID: 154359) (JIRA ID: EM-23715)
- Config Push processes will no longer write DEBUG logging to `silo.log` when the processes are not set to DEBUG. (Support ID: 154812) (Case: 00098404) (JIRA ID: EM-24255)

## PowerShell

- Fixed an Unhandled Exception that would occur when a PowerShell Dynamic Application's discovery object is disabled on a device. (JIRA ID: EM-30732) (Case 00034944) (Support ID: 178685)
- Addressed an issue with PowerShell credentials. Set the credential cache file to use a user-specific name (`/tmp/krb5cc_AD_domain_AD_username`). This allows multiple users to run the same Dynamic Application in the same domain without generating the error "authGSSClientInit() failed." (Support ID: 178992), (Case: 00037699), (Case: 00040151), (Case: 00037699), (Case: 00053678), (Case: 00058763), (Case: 00067455) (JIRA ID: EM-31237) (JIRA ID: EM-35061)

## Reports

- Addressed an issue where the File System Report contained entries where drives that were removed from devices were included in the report.
- Addressed an issue with a gap in data that caused unexpected billing data in the Interface Billing report. (Case: 00024606) (JIRA ID: EM-29843)
- Addressed an issue with an inconsistent number of instance of units in the Event Detections report. (JIRA ID: EM-317131) (Case: 00043546)

- Addressed an issue that occurred when generating an ad-hoc report with an Output Format of "Web page" and the URL for the report used http instead of https, which prevented the report from being displayed. (Case: Case 00024606) (JIRA ID: EM-33718)

### ScienceLogic API

- Addressed an issue that was causing a System internal error to display, rather than a 404 Not Found error, when a user tried to view a non-existing ticket in the REST API.

### Security

- Addressed a security issue for the Agent running on Linux.
- The DB password specified in /etc/silo.conf is no longer plaintext. SL1 also added introduces a script to be able to encrypt/decrypt the silo.conf DB password on demand. (Case: 00034510) (Case: 00086279) (JIRA ID: EM-25290)

### System Update

- Addressed an issue with System Update. (Case: 00081236) (JIRA ID: EM-35167). Made performance improvements to ease possible buffer problems that cause processes to hang while writing to stderr.

### Ticketing

- An issue was addressed in the Ticket Editor that was causing HTML text to appear instead of the appropriate template when users attempted to use a template in the Notes and Attachments section.

### Topology Maps

- Made performance improvements to legacy Layer-2 topology to decrease the likelihood of SIGTERMs or memory issues. In addition, if Layer-2 topology crunch fails, it will no longer fail block topology crunch for other types of topologies.
- Improved LLDP topology and relationship detection.

### Traps

- Addressed an issue inbound traps. Now, when SL1 receives a trap and that trap includes a varbind OID for MAC address, the Event Engine will evaluate the varbind OID and then the OID one level up for a match. (Case: 00080899) (Case: 00056511) (JIRA ID: EM-33158)

### User Interface

- A misspelling was corrected in the Link-Message text of the "Poller: File System usage now below (critical) threshold" Event Policy.

### Vanished Devices

- On the Vanished Device Manager page, the Filter-While-You-Type feature now works for the Hours Until Purge column. (Support ID: 110224) (JIRA ID: EM-15604)
- You can now manually purge vanished devices from the Vanished Device Registry page. (Case: 00089588) (Case: 00101009), (JIRA ID: EM-36548)



## Virtual Collector Groups

- You can now use device templates to change the alignment of virtual collector groups.

---

## PowerPacks in 10.1.0

Before upgrading to 10.1.0 , 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.

The 10.1.0 release includes the following PowerPacks that are new or updated and included with the release:

- Amazon Web Services, v116
- Cisco: Meeting Server, v101
- Cisco: IPSLA, v102
- EM7 Core Reports, v111
- F5 BIG-IP, v103
- Microsoft: Azure, v111
- Microsoft: Windows Server, v108
- Nutanix, v104
- Supplemental Device Class Pack, v102
- Tandberg: Infrastructure, v108
- Video Reports, 102
- VMware v214

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

**CAUTION:** If you are currently using the Amazon Web Services PowerPack included in the 8.1.0 platform release or included in an earlier platform release, please read the release notes for Amazon Web Services before migrating to later versions Amazon Web Services PowerPack. You must first migrate to Amazon Web Services, v100 before upgrading to later versions of the Amazon Web Services.

**CAUTION:** If you are currently using the VMware: vSphere Base Pack, v202 or earlier, please read the migration instructions in the release notes for VMware: vSphere Base Pack, v203 and migrate to v203 before upgrading to later versions of the PowerPack.

**NOTE:** The "Video Performance App" PowerPack is no longer included with platform releases.

**NOTE:** The "Polycom Infrastructure" PowerPack is no longer included with platform releases.

**NOTE:** The "Microsoft SharePoint" PowerPack is no longer included with platform releases. The PowerPack is superseded by the "Microsoft: SharePoint Server" PowerPack.

**NOTE:** The "Microsoft SQL Server" PowerPack is no longer included with platform releases. The PowerPack is superseded by the "Microsoft: SQL Server" and "Microsoft: SQL Server Enhanced" PowerPacks.

**NOTE:** The "Microsoft Exchange" PowerPack is no longer included with platform releases. The PowerPack is superseded by the "Microsoft: Exchange Server 2010" and "Microsoft: Exchange Server 2013" PowerPacks.

**NOTE:** The "Microsoft HyperV" PowerPack is no longer included with platform releases. The PowerPack is superseded by the PowerPack "Microsoft: Hyper-V Server".

---

## New and Updated Packages in 10.1.0

10.1.0 includes multiple package updates to address security issues and improve performance. These include packages for the Oracle Linux and for Linbit.

**WARNING:** *10.1.0 includes important security updates that require you to reboot all AIO appliances and all appliances in the Distributed stack after installing 10.1.0. If you would like*

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

---

## Beta Features

The following beta features are included in the 10.1.0 release:

### Agent

- You can now install the SL1 agent on a Solaris server. This is a beta feature in 10.1.0.
- You can now install the SL1 agent on an AIX server. This is a beta feature in 10.1.0.

### Anomaly Detection

- SL1 includes a new beta feature, anomaly detection. Anomaly detection can examine CPU data and all data from Dynamic Applications to and detect anomalies for devices, including devices that are included in Business Services.

### Publisher

- SL1 includes a new beta feature, Publisher. Publisher allows users to consume data collected by SL1 (via Data Collectors and the ScienceLogic Agent) in real-time, outside the user interface. Publisher can provide data for long-term storage or provide input to other applications the perform analysis or reporting.

### SNMP Collection

- Added a new beta 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.
- Users can collect metrics for the new beta feature, Concurrent SNMP Collection.

### PowerPacks

- For a list of beta versions of PowerPacks, go to the [ScienceLogic Support Site](#) and search for "Beta".

---

## 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 silo.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/silo.conf`. In the LOCAL section, add the line:  
`kbase_disabled=1`
3. Use an editor like vi and edit the file `/etc/siteconfig/siloconf.siteconfig`. In the LOCAL section, add the line:  
`kbase_disabled=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 10.1.5.1.

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 firewallld **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 Big Ben 8.12.0

EditDelete

Release Version Name  
SL1 Big Ben 8.12.0

End of Maintenance  
10/31/2020

Online Documentation

Number of Files  
2

Version  
8.12.0

End of Life  
4/30/2021

Allow Customers to View on Community  
☐


Release Files (2)

New

FILE NAME	COMMENTS	RECORD TYPE	RELEASE DATE
<a href="#">8.12.0</a>	System running 8.10.0+ or System has 8.10.0 siloupdat...	Product Update	4/24/2019
<a href="#">8.12.0</a>		Image	4/24/2019

View All

- 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.

## Upgrade Process for SL1 Extended Architecture

For upgrade details for SL1 Extended Architecture, see the chapter on *Upgrading SL1* in the [System Administration](#) manual or view that chapter [online](#).

For instructions on how to deploy the SL1 Extended Architecture, see the [Installation](#) manual or view that chapter [online](#).

## Special Steps for Upgrading MariaDB

SL1 10.1.x includes an upgrade to MariaDB. The upgrade did not include a tool, jemalloc, that helps manage memory usage.

**NOTE:** For SL1 versions 10.2.0 and later, jemalloc is included with the platform. For SL1 versions prior to 10.1.0, jemalooc is included with the platform.

To avoid problems with memory usage on Database Servers, perform the following steps after upgrading MariaDB for 10.1.x.

**NOTE:** Perform these steps first on the active Database Server and then on each additional Database Server in your SL1 system.

1. Open an SSH session to the Database Server.
2. To verify that the Database Server is not currently running jemalloc, enter the following at the shell prompt:  

```
silo_mysql -e 'show global variables like "version_malloc_library"
```

  
If the Database Server is not currently running jemalloc, the shell will display the following:

Variable Name	Value
version_malloc_library	system

3. Search for the file `/usr/lib64/libjemalloc.so.1`.

If the file does not exist, contact ScienceLogic Customer Support to request the file `jemalloc-3.6.0-1.el7.x86_64.rpm`.

To install the RPM, use a file-transfer utility, copy the file to a directory on the SL1 appliance. Then enter the following at the shell prompt:

```
cd /usr/lib64
sudo yum install jemalloc-3.6.0-1.el7.x86_64.rpm
```

4. Create the file `/etc/systemd/system/mariadb.service.d/jemalloc.conf`:

```
vi /etc/etc/systemd/system/mariadb.service.d/jemalloc.conf
```

5. Add the following lines to the file:

```
[Service]
Environment="LD_PRELOAD=/usr/lib64/libjemalloc.so.1"
```

6. Save and close the file.

7. Reload the systemd config files:

```
sudo systemctl daemon-reload
```

8. Restart the Database Server:

To restart the **standalone Database Server** or the **primary Database Server in a cluster**, enter the following:

```
sudo systemctl restart mariadb
```

To restart each **secondary Database Server in a cluster**:

- a. Open an SSH session to the secondary Database Server. At the shell prompt, enter:

```
coro_config
```

- b. Select **1**.

- c. When prompted to put the Database Server into maintenance, select **y**.

- d. Open an SSH session to the primary Database Server. To pause SL1, enter the following at the shell prompt:

```
sudo touch /etm/.proc_mgr_pause
```

- e. In the SSH session for the secondary Database Server, restart MariaDB:

```
crm resource restart mysql
```

- f. After MariaDB has restarted successfully on the secondary Database Server, return to the SSH session on the primary Database Server. Remove the pause file for SL1:

```
sudo rm /tmp/.proc_mgr_pause
```

- g. In the SSH session on the secondary Database Server, take the Database Server out of maintenance. At the shell prompt, enter:

```
coro_config
```

- h. Select **1**.

- i. When prompted to take the Database Server out of maintenance, select **y**.

9. To verify that jemalloc is running on the Database Server, enter the following at the shell prompt:

```
silo_mysql -e 'show global variables like "version_malloc_library"'
```

If the Database Server is currently running jemalloc, the shell will display something like the following:

Variable Name	Value
version_malloc_library	jemalloc 3.6.0-0-g46c0af68bd248b04df75e4f92d5fb804c3d75340

10. Perform these steps on each Database Server in your SL1 system.

---

## Manually Updating ap2

If you use the unified user interface (ap2) and have previously updated the user interface using the Content Management page, you must perform these steps after upgrading to 10.1.0. This is a one-time manual process. After performing these steps, future releases will not require this step.

If you have previously updated ap2 using the Content Management page, you might see the following errors after upgrading to 10.1.0:

- URLs that include "/ap2" return a 502 error
- URLs that include "/gql" return a 502 error
- The new HTML5-based maps are not working in /ap2

To resolve these errors, perform the following steps on **each All-In-One Appliance, Administration Portal and Database Server** in your SL1 system:

1. Either go to the console of the All-in-One, Administration Portal, or Database Server or use SSH to access the All-In-One, Administration Portal, or Database Server.
2. Open a shell session on the server.
3. Determine the version of /ap2 that is installed in the /opt partition. To do this, enter the following at the shell prompt:

```
cat /opt/em7/nextui/node_modules/@sciencelogic/ap2/package.json | grep version
```



4. You should see something like:

```
"version": "5.222.1-release.25"
```

Note the version number.

5. Determine the version of /ap2 that is installed in the /var partition. To do this, enter the following at the shell prompt:

```
cat /var/opt/em7/nextui/lib/node_modules/@sciencelogic/ap2/package.json | grep  
version
```

6. You should see something like:

```
"version": "5.1555.0-hotfix.8"
```

Note the version number.

7. The version number in the /opt partition should be later than the version number in the /var/partition.
8. To prevent SL1 from trying to load the older version of ap2 in the /var partition, rename the directory. To do this, enter the following at the shell prompt:

```
mv /var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2  
/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2_old_<today's date>
```

For example:

```
mv /var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2  
/var/opt/em7/gui/nextui/lib/node_modules/@sciencelogic/ap2_old_20200825
```

9. Restart the nextUI service. To do this, enter the following at the shell prompt:

```
sudo systemctl restart nextui
```

---

## 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)

---

## Future Deprecations

The following features will be deprecated in 10.2.0:

- The IPMI features(EM-38111)
- View > Other Views (EM-38109)
- Views > Topology Maps will be replaced by newer mapping features for the topology maps (EM-38108)
- Views >Classic Views will be deprecated (EM-38107)
- The following view will be deprecated: (EM-38106):
  - Views >Device Maps >Device Groups
  - Views >Device Maps >Organizational
  - Views > Device Maps >Device Categories

- System Usage Pie Chart (EM-38105)
- All Flash graphics will be removed from SL1. (EM-37894)
- The Event Overview will be replaced by a tabular event overview. (EM-37893)
- The Hardware Inventory page (EM-37892)
- The Org Clock EM-37890)
- The Inbox tab EM-37888)

---

## Known Issues

To view the list of Known Issues and workarounds, see the ScienceLogic Customer Portal at <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](mailto:legal@sciencelogic.com)



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

International: +1-703-354-1010