



SL1 11.2.4.1 Release Notes

SL1 version 11.2.4.1 (Document revision 3)

Overview

SL1 11.2.4.1 includes multiple package updates to improve security and system performance.

This document covers the following topics:

<i>Upgrade Notes</i>	3
<i>Caveats</i>	11
<i>Known Issues for 11.2.4.1</i>	12
<i>Issues Addressed in 11.2.4.1</i>	13
<i>New Features in 11.2.4</i>	13
<i>Issues Addressed in 11.2.4</i>	14
<i>New Features in 11.2.3</i>	14
<i>Issues Addressed in 11.2.3</i>	15
<i>New Features in 11.2.2</i>	16
<i>Issues Addressed in 11.2.2</i>	17
<i>New Features in 11.2.0</i>	18
<i>Issues Addressed in 11.2.0</i>	30
<i>SL1 Extended Architecture</i>	40
<i>PowerPacks in 11.2.4.1</i>	41
<i>Upgrade Process for Systems Running 8.1.0 and Earlier</i>	43
<i>Upgrade Process for Systems Running 8.1.1 and Later</i>	43
<i>Recently Deprecated Features</i>	44

Upgrade Notes

WARNING: All customers must upgrade to SL1 version 12.1.1 or later and convert to OL8 by October 31, 2024, or before upgrading to SL1 version 12.2.0. **If you take no action, all older SL1 systems with OL7 will continue to run, but ScienceLogic will not support them and the systems might not be secure.**

If you plan to upgrade to SL1 12.1.1, you **should not** consume version 11.2.4.1; there is no supported upgrade path from 11.2.4.1 to 12.1.1. Instead, you should upgrade to 11.2.4, from which you can then upgrade to 12.1.1.

WARNING: There is a known issue that is preventing some users from logging in to the user interface after upgrading from an SL1 11.x release to version 11.2.4.1, due to the system reporting that your SL1 appliance is not licensed. This is happening in two specific upgrade scenarios:

- New customers who have never had a permanent license and are upgrading from the 11.2.0 ISO to 11.2.4.1.
- Existing customers who previously upgraded from an earlier release to 11.2.0 and are now upgrading to 11.2.4.1 and whose original license expiry date is prior to the date they upgrade to 11.2.4.1.

To work around this issue, prior to upgrading to version 11.2.4.1, you can run a Python script on your active Database Appliance that will check for licenses that are about to expire in a predefined amount of time. You can then proactively update any expired or soon-to-expire licenses prior to upgrading to SL1 11.2.4.1.

You can download this pre-upgrade script by clicking the following link: https://docs.sciencelogic.com/release_notes/check_license_expired.py.

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

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

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

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

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

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

CAUTION: Due to the upgrade to MariaDB 10.4.28 that is required to consume SL1 11.2.4.1, you can no longer execute multi-statement SQL operations in Python in this release. This issue should only impact snippets and Dynamic Applications that use custom code with statements like `"dbc.execute("SELECT 1; SELECT 2; SELECT 3")"`, which was a practice that was already discouraged due to the excess load such statements can cause on the Database Server.

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

This change significantly decreases the number of Microsoft Windows servers that can be supported on each Data Collector in your SL1 system. Users who need to monitor Windows devices using WMI should analyze their system resources and capacity before upgrading to 11.2.4.1. For guidance about sizing, see the updated [Collector Sizing guidelines for WMI endpoints](#).

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

CAUTION: If you are using one of the following PowerPacks, you must upgrade to the specified version before upgrading to SL1 version 11.2.4.1:

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

- NetApp: Base Pack v106
- Oracle: Database v104
- Oracle: MySQL v102
- SL1: Concurrent PowerShell Monitor v103
- VMware Automation v102
- Windows PowerShell Automation v104

The previous versions of these PowerPacks will not prevent SL1 version 11.2.4.1 from installing or operating, but they might not collect data post-upgrade due to technical incompatibilities.

CAUTION: If you are using one of the following PowerPacks, you must edit an HTTP header in the credential before you upgrade:

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

To edit the credential HTTP header:

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

CAUTION: If you are using Aurora RDS (MySQL 5.7) with SL1, you must update to the following PowerPack before installing the 11.2.4.1 release:

- Linux Base Pack v105

If you are using Aurora RDS (MySQL 5.7) with SL1, the *Cisco: Cloud Services Platform* PowerPack and the *Cisco: UC VOS Applications* PowerPacks will not work with the 11.2.4.1 release until the following upgraded PowerPacks are available:

- Cisco: Cloud Services Platform v107
- Cisco: UC VOS Applications v110

If you are using Aurora RDS (MySQL 5.6) with SL1, older versions of the *Linux Base Pack*, *Cisco: Cloud Services Platform*, and *Cisco: UC VOS Applications PowerPacks* will continue to work with the 11.2.4.1 release.

NOTE: Version 103 of the *ScienceLogic Support Pack PowerPack* replaced the "Support: MySQL Performance" Dynamic Application with the "Support: MariaDB Performance" Dynamic Application.

- If you are using AWS Aurora RDS with SL1, you cannot upgrade from the following SL1 versions to 11.2.4.1 due to an Aurora RDS compatibility issue:
 - 10.2.5
 - 10.2.6
 - 10.2.6.1
 - 10.2.7
 - 11.1.3
 - 11.1.4
 - 11.1.5
 - 11.1.6

This issue does not impact other deployment types, and is planned to be resolved in an upcoming release.

If your SL1 system uses Amazon RDS for storage and you plan to upgrade to SL1 11.2.4.1 from a version prior to 11.2.0, then immediately after upgrading SL1, you must also upgrade the MySQL database to version 5.7 with an Aurora MySQL engine of 2.x.

If you upgrade to SL1 11.2.4.1 but do not immediately upgrade the MySQL database to version 5.7 with an Aurora MySQL engine of 2.x, then you will experience an outage and the SL1 user interface will be unavailable until MySQL is upgraded.

To plan an upgrade strategy that minimizes downtime, ScienceLogic strongly suggests that you contact your Customer Success Manager before upgrading to SL1 11.2.4.1 and upgrading Amazon RDS (Aurora MySQL).

- After installing 11.2.4.1, you can upgrade to the 11.3.1 or 12.1.0 releases regardless of your deployment type.
- If you are upgrading from a version of SL1 prior to 11.1.0, during an upgrade to 11.2.4.1, the user interface will be unavailable due to an upgrade from PHP version 5 to PHP version 7, which was introduced in SL1 version 11.1.0.

- **If you are running a version of SL1 prior to 8.12.0, do not install version 11.2.4.1 if you do not plan to immediately consume version 11.2.4.1.** After you import the 11.2.4.1 release, all appliances in your SL1 system will now use the **new system update**. After you import the 11.2.4.1 release, you will not be able to stage and deploy any versions of SL1 prior to 8.12.0 or apply patches to any versions prior to 8.12.0. For details on the new system update, see the [release notes for SL1 version 8.12.0](#).
- **To install 11.2.4.1, you must have already imported, staged, and deployed SL1 version 8.6.0 or a later release.**
- SL1 version 11.2.4.1 includes important security updates. **You must reboot all All-in-One appliances and all appliances in the SL1 Distributed Architecture after deploying version 11.2.4.1.** 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](#) for SL1 (<https://support.sciencelogic.com/s/known-issues#sort=relevancy>) before installing a new update. For Known Issues specific to this release, see the [Known Issues](#) section of 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.
- Some SL1 updates require you to upgrade MariaDB and then reboot all SL1 appliances to apply important security updates.

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

SL1 Release	Required MariaDB Version	Requires Reboot?
11.2.4.1	10.4.28	Yes
11.2.4	10.4.28	Yes
11.2.3.1	10.4.28	Yes
11.2.3	10.4.28	Yes
11.2.2	10.4.26	Yes
11.2.0	10.4.24	Yes
11.1.6.1	10.4.28	Yes
11.1.6	10.4.28	Yes
11.1.5	10.4.26	Yes
11.1.4	10.4.26	Yes
11.1.3	10.4.25	Yes
11.1.2	10.4.24	Yes
11.1.1	10.4.22	Yes
11.1.0.1	10.4.20	Yes
11.1.0	10.4.20	Yes
10.2.7	10.4.27	Yes
10.2.6.1	10.4.26	Yes

SL1 Release	Required MariaDB Version	Requires Reboot?
10.2.6	10.4.26	Yes
10.2.5	10.4.22	Yes
10.2.4.1	10.4.22	Yes
10.2.4	10.4.22	Yes
10.2.3	10.4.21	Yes
10.2.2	10.4.18	Yes
10.2.1	10.4.18	Yes
10.2.0	10.4.18	Yes

NOTE: For instructions on updating MariaDB or rebooting SL1, see the chapter on "Upgrading SL1" in the [System Administration](#) manual.

Upgrading PhoneHome Collectors

NOTE: If you do not currently use a PhoneHome configuration, then you can skip this section.

The system update procedure includes a new pre-upgrade test that checks for existing PhoneHome Database Servers. This test looks for PhoneHome token IDs inside the `/home/phonehome0/config.json` file and fails if the value of the token ID field is less than or equal to "0".

In previous versions of SL1, the primary PhoneHome Database was not self-registered with a token, causing it to have an ID of "0". Therefore, if you have an existing PhoneHome configuration, then you must perform these one-time manual configuration steps on all Database Servers in your PhoneHome configuration prior to upgrading to SL1 version 11.2.4.1:

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

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

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

```
phonehome status
```

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

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

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

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


6. Repeat steps 3-5 for all PhoneHome Database Servers that have an ID value of "0".
7. Type the following command to ensure that all of your PhoneHome Database Servers are synced:

```
phonehome sync
```

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

NOTE: Do not attempt to upgrade to 11.2.4.1 until all pre-upgrade tests are successful on all PhoneHome Database Servers.

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

LDAP Authentication

This section describes the various LDAP authentication configurations that are supported in SL1.

CAUTION: If you are using an LDAP configuration other than one that is listed below, you should contact ScienceLogic Support or your Customer Success Manager to explain your use case. Non-supported configurations will be deprecated in a future release.

Configuration 1: Basic LDAP Authentication

- Configure an authentication profile that lists *EM7 Login Page* as the aligned credential source.
- The aligned authentication resource is associated with an LDAP/AD credential.
- It does not matter if the aligned LDAP/AD credential uses the %u or %e variables in its RDN string or if the RDN string is a defined value. If it is a defined value, it must also include a bind password.
- Optionally, an SSL certificate has been imported and installed if you are using LDAP over SSL.

NOTE: You can log in through REST API using an LDAP configuration.

Configuration 2: LDAP Configuration for CAC Authentication

- Configure one authentication profile, for most uses:
 - The authentication profile lists *CAC/Client Cert* as the aligned credential source.
 - The aligned authentication resource is associated with an LDAP/AD credential.
 - The aligned LDAP/AD credential uses a defined RDN string with a bind password; it cannot use the %u or %e variables in its RDN string.

- Configure a second authentication profile for administrator or maintenance access:
 - The authentication profile lists *EM7 Login Page* as the aligned credential source.
 - The aligned authentication resource is the *EM7 Internal* resource.

NOTE: You cannot log in through REST API using CAC authentication.

NOTE: You cannot have both CAC and non-CAC LDAP users on the same SL1 system.

NOTE: To disable a user's CAC authentication access, remove the user from the LDAP/AD server.

Configuration 3: Multiple LDAP Authentication Resources Used in the Same Authentication Profile

- Configure an authentication profile that lists *EM7 Login Page* as the aligned credential source.
- The authentication profile lists multiple aligned authentication resources, all of which are associated with LDAP/AD credentials.
- It does not matter if the aligned LDAP/AD credentials use the %u or %e variables in their RDN strings or if the RDN strings are a defined value. If they are defined values, they must also include bind passwords.
- Optionally, an SSL certificate has been imported and installed if you are using LDAP over SSL.

Configuration 4: One LDAP Authentication Resource Used in Multiple Authentication Profiles

- Configure one authentication profile:
 - The authentication profile lists *EM7 Login Page* as the aligned credential source.
 - The aligned authentication resource is associated with an LDAP/AD credential.
 - It does not matter if the aligned LDAP/AD credential uses the %u or %e variables in its RDN string or if the RDN string is a defined value. If it is a defined value, it must also include a bind password.
 - Optionally, an SSL certificate has been imported and installed if you are using LDAP over SSL.
- Configure a second authentication profile:
 - The authentication profile lists *EM7 Login Page* as the aligned credential source.
 - The aligned authentication resource is same one used in the first authentication profile.

Configuration 5: Basic HTTP Authentication with LDAP

- Configure an authentication profile that lists *HTTP Auth* as the aligned credential source.
- The aligned authentication resource is associated with an LDAP/AD credential.
- It does not matter if the aligned LDAP/AD credential uses the %u or %e variables in its RDN string or if the

RDN string is a defined value. If it is a defined value, it must also include a bind password.

- Optionally, an SSL certificate has been imported and installed if you are using LDAP over SSL.

Caveats

Consider the following caveats before deploying 11.2.4.1:

- In SL1 version 11.1.0, all PHP code was converted to PHP 7. If you have created custom content in PHP, see this page for backward compatibility: <https://www.php.net/manual/en/migration70.incompatible.php>
 - Due to the PHP upgrade in version 11.1.0, versions of Global Manager prior to 11.1.0 will not work with SL1 11.2.4.1.
 - Due to the PHP upgrade in version 11.1.0, Web Proxy Services will not work in SL1 11.2.4.1.
 - PowerPacks built in SL1 version 11.1.0 and later releases cannot be imported into previous versions of SL1 due to the upgrade to PHP version 7.0. However, PowerPacks built in previous releases of SL1 can be imported into version 11.2.4.1.
- As of version 10.1.0, SL1 no longer includes Flash.
- As of version 8.12.2, ScienceLogic no longer updates the help that appears when you click the **[Guide]** button that appears in the classic user interface. The Unified UI provides a new tool for inline help. Clicking the **[Help]** button at the top of the page now opens a Help topic about that page on the right-hand side of the page. From that topic, you can also click a context-sensitive link to open the relevant page in the product documentation at docs.sciencelogic.com in a new browser window.
- As of 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. (For more details, see <https://support.sciencelogic.com/s/article/1161> and <https://support.sciencelogic.com/s/article/1423>.) ScienceLogic recommends that you back up this file before applying an update and then reapply your changes to this file.
- SL1 8.10.0 and later releases do not support Data Collectors and Message Collectors running the CentOS operating system. **If your system includes Data Collectors and Message Collectors running the CentOS operating system, contact your Customer Success Manager for details on upgrading Data Collectors and Message Collectors to Oracle Linux before installing 11.2.4.1.**
- SL1 Extended Architecture does not provide MUD support.
- During the normal system update process, multiple processes are stopped and restarted. This might result in missed polls, gaps in data, and/or unexpected errors. ScienceLogic recommends that you always install ScienceLogic releases during a maintenance window.

- The ScienceLogic system update process starts a background process that can perform any required post-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.

Known Issues for 11.2.4.1

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

The following known issues exist for SL1 version 11.2.4.1:

- After upgrading to 11.2.4.1, if you attempt to generate a one-time password for the sl1admin user, the resulting password might appear encrypted or like a random string of characters. In this case, you might need to rebuild hashes for one-time passwords. You can fix this by logging in to the console of the primary Database Appliance and running the following MySQL command: (Jira ID: EM-51394)

```
update master_platform.appliance_access_hash set disabled = 1;
```

- With on-premises SL1 Extended systems, the TLS handshake can fail between the Windows Agent on a monitored device and the SL1 streamer service. (Jira ID: DO-4079) (Jira ID: DO-4115). For details and the workaround, see: <https://support.sciencelogic.com/s/article/7417>
- If you use CAC with LDAP/AD, then to configure CAC, users must define an LDAP service account with permissions that allow the service account to query LDAP. (Case: 00207174) (Jira ID: EM-46927)
- In Global Manager systems, when the Global View button is toggled on, SL1 returns only partial data if you use the filtering feature at the top of the Devices page. Additionally, an error message appears in the Recent Devices panel of the Activity Center. (Jira ID: SLUI-14346)
- When creating or updating credentials in the AP2 user interface, SL1 fails to save the information you entered in the credential fields if you move to a new field too quickly or if you attempt to save the credential too quickly. If you experience this issue, ScienceLogic recommends editing the credential again, but using a slower pace when moving between fields and when saving. Alternatively, you can create or update some credential types in the classic SL1 user interface. (Jira ID: EM-46350)
- Users who deploy SL1 on AWS in a high-availability configuration might experience an issue where they cannot import PowerPacks on their active node. If this occurs, you can log in to the classic SL1 user interface on passive node to import the PowerPack.
- In new installations, some PowerPacks that are normally installed by default are not being installed. This behavior has been observed with the "EM7 Web Server," "F5 Big-IP," "Microsoft Azure," "Microsoft Base Pack," "Microsoft: Windows Server," "SL1 Default Dashboards," and "Supplemental Device Class" PowerPacks. You can manually install these PowerPacks after SL1 has been installed and configured. For instructions, see the section on [Installing a PowerPack](#) in the **PowerPacks** manual. This issue does not impact SL1 instances that have been upgraded from earlier releases. (Jira ID: SOL-25047)

Upgrading to SL1 11.2.4.1 with AWS Aurora

The user interface in SL1 11.2.4.1 is not compatible with versions of the AWS Aurora database supported by versions of SL1 prior to 11.2.0.

If your SL1 system uses AWS Aurora for storage and you plan to upgrade to SL1 11.2.4.1 from a version prior to 11.2.0:

- You must also plan to upgrade the AWS Aurora database immediately after upgrading SL1.
- You must upgrade the Aurora database to RDS 5.7 with engine 2.07.6 or later.
- After updating to SL1 11.2.4.1, the user interface will be unavailable until AWS Aurora is updated.

If you upgrade to SL1 11.2.4.1 but do not immediately upgrade the AWS Aurora database to RDS 5.7 with engine 2.07.6 or later, you will experience an outage.

To plan an upgrade strategy that minimizes downtime, ScienceLogic strongly suggests that you contact your Customer Success Manager before upgrading to SL1 11.2.4.1 and upgrading AWS Aurora.

Issues Addressed in 11.2.4.1

Security

- SL1 11.2.4.1 includes multiple package updates to improve security and system performance. Among other things, these updates address the following "high" and "moderate" security vulnerabilities: CVE-2018-7160, CVE-2020-15862, CVE-2022-29224, CVE-2022-29225, CVE-2022-29226, CVE-2022-29227, CVE-2022-292248, CVE-2022-31045, CVE-2022-32212, CVE-2022-37434, CVE-2023-0286, CVE-2023-0767, CVE-2023-2650, CVE-2023-2975, CVE-2023-3446, CVE-2023-37276, CVE-2023-37920, CVE-38325, CVE-2023-38408, CVE-2023-38545. (Jira IDs: EM-60209, EM-60342, EM-60364, EM-60367, EM-60477, EM-60496, EM-60507, EM-61783, EM-61816, EM-62181, EM-62246, EM-62389, DO-5611, CPL-644, CPL-645)

New Features in 11.2.4

Data Collection

- With this release, the SSH Collector container was removed from SL1. To support this change, the "Data Collection: SSH Collector" process is no longer available in new installations of SL1.

IMPORTANT: If you are upgrading from an earlier release and you were previously using the SSH Collector, you must reboot any Data Collectors that were previously using the "Data Collection: SSH Collector" process. After the upgrade, you can go to the **Appliance Manager** page (System > Settings > Appliances) to determine which appliances might require a reboot.

Operating System

- A notice now appears when accessing an SL1 stack using the command line interface indicating the end-of-life date for Oracle Linux 7 and providing guidance for users about the necessary upgrade paths.

Security

- 11.2.4 includes multiple package updates to improve security and system performance. Among other things, these updates address the following "high" and "moderate" security vulnerabilities: ELSA-2023-1332, ELSA-2023-1335, ELSA-2017-2492, ELSA-2023-1095, ELSA-2023-3555, and ELSA-2023-3556.

Issues Addressed in 11.2.4

Data Collection

- Addressed an issue in which SNMP collections were missed by the Concurrent SNMP Collection Service during unexpected service restart. (Case: 00288208) (Jira ID: EM-59862)
- Addressed an issue that was causing average interface utilization metrics to exceed 100%. (Cases: 00318619, 00351240) (Jira ID: EM-57005)

New Features in 11.2.3

API

- Added two new fields to the SLAmin API's `status` resource to indicate whether your SL1 system is using a "classic" PhoneHome configuration from SL1 version 11.1.x and earlier or if it is using the "new" PhoneHome configuration from version 11.2.0 and later. These fields are:
 - `is_new_phonehome`, which returns a boolean value indicating if you are using the new PhoneHome configuration.
 - `phonehome_rpm_version`, which returns a string that represents the PhoneHome RPM version.

Deployment

- SL1 11.2.3 includes an upgrade to Kubernetes 1.24.

Device Groups

- Made numerous performance improvements in the Device Group membership evaluation process.

Devices

- Made several updates to device merge and vanish functionality to ensure those features affect credentials and device relationships as intended.

PhoneHome

- In SL1 11.2.3, Military Unique Deployment (MUD) users are able to configure PhoneHome communications using the Web Configuration Utility that was previously used for that function prior to SL1 11.2.0. Non-MUD users must continue to use the Node Configuration Utility that was introduced in 11.2.0.

Security

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

Issues Addressed in 11.2.3

Classic IT Services

- Updated classic IT services to ensure that dynamic rule matches are recalculated whenever IT service changes are saved. (Jira IDs: EM-56014, EM-55303) (Case: 00311745)

Data Collection

- Addressed an issue that was causing unhandled exceptions and critical system events related to file system statistics to occur during data collection. (Jira IDs: EM-55603, EM-55141, EM-55870, EM-56778, EM56774, EM-56390, EM-57223, EM-57120, EM-55529) (Case: 00297434) (Case: 00308269) (Case: 00328596)

Discovery

- Invalid device interface IP addresses are now filtered out during discovery to ensure the discovery process completes. (Jira ID: EM-53177) (Case: 00267240) (Case: 00278024) (Case: 00286797)

Dynamic Applications

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

Events

- Addressed an issue that was resulting in event filtering not working as intended. (Jira IDs: EM-57316, EM-57231) (Case: 00320869) (Case: 00330683)

PhoneHome

- Addressed an issue that was causing the phd (PhoneHome server) service on Data Engines to fail for SaaS/Amazon RDS environments. With this change, the SL1 11.2.3 version of PhoneHome requires Amazon Aurora MySQL version 2 for SaaS/Amazon RDS environments. If you do not have version 2 installed, the phd (phonehome server) service on Data Engines fail with the following error: "Error: tls: server selected unsupported protocol version 301". (Jira ID: EM-56128)(Case: 00327348)

User Accounts

- Resolved an issue in which non-administrator users were unable to access all organizations and devices due to a "GraphQL Error: Forbidden" error. (Case: 00319376) (Jira ID: SLUI-16424)

User Interface

- Updated the **[Registered]** tab of the **Nodes** page to increase the number of appliances that display to 25 and to automatically display the next 25 appliances when you scroll to the bottom of the list. (Jira ID: EM-56170) (Case: 00319543) (Case: 00320829)

New Features in 11.2.2

Licensing

- As of 11.2.2, the default SL1 auto-license expiry is set to 90 days for ISO installations, and the default device capacity in the auto-generated license is 1,000 devices.

IMPORTANT: In SL1 version 11.2.0, the **Licensing** pages were removed from the Web Configuration Utility. Data Collector, Message Collector, and Administration Portal appliances are automatically licensed when they are added to your SL1 stack. However, Database Server and All-In-One Appliances must still be licensed before you can use this release. To do so, you will need to submit a license request to ScienceLogic Support. For more information, see Knowledge Base article: <https://support.sciencelogic.com/s/article/9617>.

Security

- 11.2.2 includes package updates that improve security and performance. Among other things, these updates address two "critical" security vulnerabilities, ELSA-2022-6834 and ELSA-2022-6878.

Issues Addressed in 11.2.2

API

- Addressed an issue in which the `/api/monitor` process was not found after upgrading to SL1 11.2.0 (Case: 00273881) (Jira ID: EM-53430)

Data Retention

- Updated the Hourly Maintenance Device Pruner to be more efficient and purge a list of devices instead of a single device. Purging the devices tasks is set as the final task to execute so other hourly maintenance tasks are finished first. (Case: 00256727)

Monitoring Policies

- Updated log file monitoring policies to ensure that they can be aligned properly using templates on agent-monitored devices. (Case: 00277733) (Jira IDs: EM-53356, EM-53603)

Platform

- Updated the name and default value of the setting used to define the wait time for Config Push result messages. The name was changed from `"message_timeout"` to `"result_wait_timeout"`. This update only affects users who have manually added the `"message_timeout"` configuration to a `"[CONFIG_PUSH]"` section of their `"/etc/silo.conf"` configuration file on their central database. Users need to update the timeout name. The default timeout value was changed from `"60"` to `"300"`. (Case: 00283693)(Jira ID: EM-53433)
- Updated the Config Push lookup table schema to support collector group ID values greater than 255. (Case: 00295890) (Jira ID: EM-54200)

PowerShell Collection

- Updated the Docker image version that is included with SL1 to address an issue that was causing concurrent PowerShell collections to fail on some Data Collectors that had more than 400 devices. (Jira ID: EM-54488)

System Upgrade

- Addressed an issue that was causing the SL1 pre-upgrade check to fail for `"sl1-vault-services"` due to an RPM checksum issue. (Jira ID: SLS-665)

New Features in 11.2.0

Agent

- The `sl_splice` library and helm charts used for the Agent Pipeline were updated.
- On SL1 systems running agents, SL1 automatically updates the value of the "Responder API Base URL" field on the Behavior Settings page (System > Settings > Behavior).
- SL1 environments that have the Collector Pipeline enabled but do not use the ScienceLogic Agent no longer publish messages with an empty list.
- If you specified "60 Seconds" for the Upload Interval in the Agent section of the Settings tab of the Device Investigator page for an agent device, the summary upload now includes "watched" or "monitored" files, just like the snapshot upload does. This feature is currently behind an AP2 feature toggle.
- Dynamic Applications that leverage the Snippet framework CLI/SSH toolkit can execute using the ScienceLogic Agent.
- The publishing of statsd metrics from backend services was disabled, as these metrics were never consumed by SL1.

Anomaly Detection (AIML)

- The Anomaly Detection column that appears on the Machine Learning page, the Machine Learning tab of the Device Investigator, and the Anomalies widget of the Service Investigator now indicates the build status for the metric that SL1 is evaluating for anomalies. The possible statuses that display in this column include Queued, Building, Failed, Enabled, and Disabled.
- If an anomaly detection prediction policy has been stuck in a queued status for some metrics for too long, SL1 will re-queue it.
- The week-hour prediction model that is used in machine learning-based anomaly detection was updated to better handle situations where the data is flat, and thus the standard deviation is zero.
- Improved anomaly detection model selector logging.
- Machine learning-based anomaly detection will now more frequently reselect the model used to detect anomalies, to better adapt to new data patterns.
- Improved the signal-to-noise ratio in machine learning-based anomaly detection and improved scalability by preloading input cache with historical data when making the first prediction with an anomaly detection model.

Backups

- The Backup Management page (System > Settings > Backup) now includes fields for Full Backup on Disaster Recovery (DR), Configuration Backup on DR, Full Retention Period, DR Retention Period, and advanced settings for each type of backup (Configuration Backup, Full Backup, and DR Backup).
- Added the ability to use AWS S3 for backups.

Business Services

- Added a new Refresh Interval field to the Info drop-down in the Business Service Investigator. This field enables users to edit the frequency with which health, availability, and risk values are updated.
- Added a 'Next Refresh' value to the Business Services Overview page to let users know the next time the health, availability, and risk values will be updated. These times are based on the user's time-zone preference.
- The Change Events widget on the Business Services Overview page was relabeled as the Changes widget. In addition, the widget now displays change data from Restorepoint, indicates emergency changes from ServiceNow, and includes several additional user experience upgrades such as the ability to filter and search for events by date and an External Ticket column that enables you to view ServiceNow tickets associated with cleared events.
- Optimized the load-time for the Business Services Overview page and the Events tab in that page.
- The frequency for the "Business Services: Service Management Engine" process, which you can view and customize on the Admin Processes page (System > Settings > Admin Processes), is now set to one (1) minute by default; this allows custom poll frequencies for individual services to work correctly. Every minute, this process will run and determine which service will need health, availability, and risk calculated based on their individual poll frequencies. During SL1 updates, all individual services poll frequencies will be reset to the default value and will run based on the new environment variable "BUSINESS_SERVICES_DEFAULT_POLL_FREQUENCY" frequency.

Collector Groups

- A new Collector Groups page (Manage > Collector Groups) was added. From this page, you can add, edit, and delete collector groups, view information about existing collector groups, and search through or sort the list of collector groups.

Collector Pipeline

- Containerized services in the Collector Pipeline and Publisher now include hardened containers to improve security.

Credentials

- Added four new "universal" credential types. From the Credentials page (Manage > Credentials), you can now create credentials for AWS Assume Role, AWS EC2, and AWS IAM guided discovery workflows, as well as S3 Backup credentials.

Custom Attributes

- The Custom Attributes page (Manage > Custom Attributes) now includes a Subscriber Count column that lists the number of subscribers for each attribute.

Dashboards

- A tree map widget was added to the charts package for use in future dashboard widget visualizations.
- The Dashboard URLs on the Dashboard page now show the Dashboard's unique identifier, which is a unique 32-character alphanumeric string. You can use these identifiers to export dashboards to PowerPacks.

NOTE: Changing the name of default/PowerPack dashboards will break the links to those dashboards after they are installed in another system. However, if you change a dashboard name and then fix the link by relinking the connected dashboard, the link will persist even if the dashboard changes name again in this system.

- Users can now add or remove Dashboards as favorites on the Dashboards page.
- The "favorite" field has been added to dashboard queries in GraphQL.
- Line Chart widgets now allow users to turn any existing Line Chart into an Area Chart. Users can choose the chart's metrics to be "not stacked", "stacked", or "100%-stacked" if the Area Chart option is enabled. When the Area Chart option is disabled, the ability to stack is not allowed.
- Users can now configure a widget title to hyperlink to another dashboard, internal AP2 link, or external URL.
- Users can now view the last time any dashboard was edited or modified. Users can also sort by or hide this column if desired.
- Users can now deploy a PowerPack containing multiple dashboards that contain links to each other. Links can be SL1-internal items as well as external links.

Deployment

- This release supports Scylla 4.5 on the SL1 Storage Nodes (SN) along with the Scylla Manager (SM).

Devices

- A Changes tab was added to the Device Investigator page that lets you view active and cleared change events from ServiceNow or Restorpoint for a specific device. This feature requires user configuration before it appears on the Device Investigator page. For more information, see the "Using the Device Investigator" chapter in the [Device Management](#) manual.
- A new Logs tab on the Device Investigator page displays all of the logs for a device.

Device Groups

- A new Dynamic Rule selector was added to Device Groups. You can now select by Software Title, which matches software title and software version.
- Performance improvements were made to dynamic Device Groups.

Discovery

- New guided discovery workflows were added for Google Cloud, Windows, SNMP, and pingable devices. In addition, these guided discovery workflows were added to the Guided Discovery Workflow Base Pack.
- New sub-workflows were added for the AWS and Windows guided discovery workflows. When you begin creating a guided discovery session for AWS or Windows, a new page appears where you can select the appropriate sub-workflow, with instructions included for each. For AWS, the previous single guided workflow has been removed; you can now select sub-workflows for AWS Assume Role, AWS EC2, and AWS IAM. For Windows, you can select sub-workflows for Windows WMI, Windows PowerShell, and the Windows SL1 Agent.
- A new Discovery Session Name field was added to all guided discovery workflows. The name you enter in this field will be used to identify the guided discovery session and any unguided discovery sessions that are created by the guided discovery session.
- SL1 now displays informational notes during the creation of guided discovery sessions to provide important information to the user.
- SL1 now displays a discovery log during the final summary step of certain guided discovery workflows.
- The Apply Device Template drop-down that is used when creating an unguided discovery session was updated to address several usability issues.
- The Discovery Session Editor page (System > Manage > Classic Discovery > create button/wrench icon) was updated so that the current selected value is highlighted in each field prior to saving. Additionally, the Reset button now behaves as expected and clears the selected fields.

Dynamic Applications

- On the Dynamic Applications Manager page (System > Manage > Dynamic Applications), the Select Actions menu now includes an option to bulk disable and bulk enable Dynamic Applications.

Email

- Outbound email can now utilize OAuth 2, PLAIN, Login, or CRAM-MD5 authentication methods and SMTPS or SMTP with TLS encryption methods.

Events

- An enhancement to the GQL query for Events Count returns results faster.
- The sl_splice library was updated from version 1.3.0 to 1.3.3 for the sl_event_storage service.
- In the Event Policy editor (Events > Event Policies > create button/select existing policy > Event Message tab, or Registry > Events > Event Manager > create button/wrench icon > Advanced tab in the classic UI), the Auto Clear section now includes a list of Available Events and a list of Selected Events. Each list includes an Add button and a Delete button.
- On the Events page and the Devices page, you can now click the Expand icon next to an item name to open a new modal with additional details about that item.

- In the classic user interface, users can add events to the Admin Warning Banner (also called the Admin Notifier) from the Event Policy Manager page (Registry > Events > Event Manager). The Select Action menu in the Event Policy Manager page includes two new options: Enable Admin Banner Warning and Disable Admin Banner Warning. The Event Policy Manager page also includes a new Admin Warning column that specifies whether or not an event will be included in the Admin Warning Banner.

Global Manager

- A new Organization table widget is now available for use in the Global Manager system dashboards.
- The Global View toggle in Global Manager can now be hidden based on user permissions. Users must have the GM_STACK_VIEW access hook align to their user policy to toggle between Global View and Local View.

GraphQL

- Added a "collectorToken" GQL resource that generates a token used to onboard new SL1 collectors.
- Added the "allocation", "endpoint", and "collectorGroup" fields to the appliance GQL resource.
- Added the "collectorConnectionRequest" GQL resource, the "collectorConnectionRequest" and "collectorConnectionRequests" queries, and the "updateCollectorConnectionRequest" mutation.
- Added the "createAppliance" GQL mutation that lets you create a new appliance or node. Please note that you cannot add an All-in-One appliance type, even if you are on an All-in-One system or a distributed system. If you add a new appliance that is a "collector" type (Message Collector or Data Collector), the status column in the database will be set to 0 (Available). If you add any other appliance type, the status column will be set to 10 (Unconfigured). The databaseUser and databasePassword can be changed only for Message Collector and Data Collector appliance types. The IP field is validated using the validator library isIP function. The name field is just an unverified string, and it does not need to be a valid IP or follow host name guidelines.
- Added the "updateAppliance" GQL mutation that lets you update an appliance or a node. You can update the IP, description, databaseUser, and databasePassword fields. The IP field is validated by checking that it is a valid IP format and also verifying the IP is not already present in the "master.system_settings_licenses" table. The databaseUser and databasePassword fields are only allowed for Message Collector or Data Collector appliance types, which is the same behavior as the create mutation.
- Refactored the appliance GQL resource from an API-driven resource to a database-driven resource. Added fields for "needsReboot", "databaseVersion", and "taskManagerPaused". The "organization" field is now deprecated, and will be removed in September 2022. Removed the previously deprecated field "updatedBy", which was renamed as "editedBy".
- Added two new GQL resources, "Schedule" and "Task", to support new scheduler features in SL1.

High Availability and Disaster Recovery

- A note was added to the [Installation](#) manual instructing users to leave port 7706 open for communication between the HA-primary and the HA-secondary and DR.

Inbox

- The Inbox feature has been deprecated and is available only for reports in the classic user interface.

IPv6

- Users can now use IPv6 addresses in addition to IPv4 when discovering devices, making queries via SQL or API, or monitoring SSL certificates or Web Content.

Installation

- The `system_status.sh` script that is run when a system is installed on VMware was reformatted to improve its readability and maintainability.
- The Rescue option on the ISO boot menu was updated to ensure that users can use the rescue interface successfully if they need to resolve any issues on their systems.
- VMware configuration checks that detect host issues that might affect VM performance were added to the `system_status.sh` script that is run when a system is installed on VMware.

Licensing and Billing

- A new user interface for viewing, adding, editing, and deleting license types was implemented.
- The Subscription Usage output (Manage>Subscription Usage) now includes visualizations of Current and Historical Calculated Usage for Standard Count licenses.
- A new "Feature Usage" tab was added to the Subscription Usage page (Manage>Subscription Usage). The new tab includes navigation content, configuration data, and usage data (time span). Feature usage information is illustrated by the following in this tab:
 - A time-series line graph and a time selector are displayed in the main pane when a user clicks the name of a Feature Performance field on the Feature Usage page. This graph displays the number of features and indexes supplied by each feature, data points associated with that feature and index over time, and a legend that shows the feature name and the index number.
 - A table graph is displayed in the main pane when a user clicks the name of a Feature Performance field on the Feature Usage page. This graph displays four columns including "Feature", "Index" (if data includes multiple indexes), the name of the Feature Configuration field, and "Storage Time". The graph's rows contain the feature name, the index, the current stored value, and the time that value was stored.
- The payload of data that is sent to ScienceLogic now includes counts of N-Tier business services and root business services.
- As part of a larger effort to overhaul the platform registration and identification, the current platform license was disabled in favor of a dummy license that enables SL1 to continue operating without needing to be reinstalled. This update should result in minimal behavior change.

Logging

- Additional audit rules were added to the audit rules file in `/etc/audit/rules.d/`.

MariaDB

- The `innodb_purge_threads` configuration setting was removed from the installation process, which now uses the MariaDB default value.
- The MariaDB slow query log is now enabled by default for all appliances that have MariaDB and rotates as needed to address any disk space issues that might occur.

Monitoring Policies

- You can now bulk enable or bulk disable monitoring policies on the Monitoring Policies pages found under Registry > Monitors. In addition, these pages now include a "State" column that displays the current state (enabled/disabled) for each policy.

Network Interfaces

- A new maintenance task was created to disable network interfaces that do not receive any traffic within a set number of days.
- A new collection process was added to poll interfaces disabled by the `'disable_inactive_interfaces'` maintenance task to check for traffic and enable collection if the interfaces start seeing traffic. The process is set to run once per day by default.

Node Configuration Utility

- The new Node Configuration Utility at `<IP address>:7700/node-config` lets you connect an SL1 Collector with an SL1 Database Server.
- The new Node Configuration Utility lets you log out and end your session when needed.

PhoneHome Collectors

- A new "phonehome" schema that stores all PhoneHome-related configuration information is now included in configuration backups.

PHP Developer Logs

- The PHP Developer Logs page (System > Tools > PHP Developer Logs) includes a new "IP in filename" checkbox. If this checkbox is selected, the name of each log file will include the IP address from which the user is logged in.
- The PHP Developer Logs page (System > Tools > PHP Developer Logs) also includes a new "Request URI" checkbox under the section on UI/REST MySQL Query Log Levels. If this checkbox is selected, the `mysqlqi.log` file will include the request URI.

Platform

- The SSL Certificate Purge Timeout threshold was added to the System Threshold Defaults page (System > Settings > Thresholds > System), the Device Template Editor (Registry > Devices > Templates > wrench icon), and the Device Thresholds (Devices > Device Manager > wrench icon > Thresholds) page.
- When an instance of the "EM7 Core: Daily Maintenance" process is terminated early because it has been running for too long, an event will be raised to alert the user of that occurrence.
- The method by which SL1 reads and accesses the silo.conf file was improved to optimize performance.

PowerPacks

- You can now apply the "Enable Selective PowerPack Field Protection" setting for individual PowerPacks on the PowerPack Properties page (System > Manage > PowerPacks > create button/wrench icon). You can still set this option globally for all PowerPacks in the Behavior Settings page (System > Settings > Behavior).
- Logging and error handling was improved in the tools for building PowerPacks.

Publisher

- Containerized services in Collector Pipeline and Publisher now include hardened containers to improve security.

Reports

- A new Device Collector Group Summary report (Reports > Run Report > Devices > Device Collector Group Summary) was added to SL1 and displays which devices are aligned with which Collector Group.
- A new Device Outage Events History report (Reports > Run Report > Devices > Device Outage Events History) was added to SL1. This report displays a list of device outages using the events table instead of the logs table, as in the Device Outage History report.
- The SL1 Core Reports PowerPack now includes the Device Collector Group Summary, Device Outage Events History, and Inactive User reports. This version also includes a number of updates to existing reports.

Run Book Automations

- A new Run Book Engine Settings page (Registry > Run Book > Run Book Engine Settings) was added that enables users to customize the settings for the Run Book Automation executor.
- Users can now trigger a restart of the Run Book Automation (RBA) Engine when they update RBA Engine settings. This restart will terminate any queued and in-progress RBA executions.
- A new access hook was added for the new RBA Executor Settings page.
- A new internal alert message will log when the Automation engine is unable to retrieve new Run Book actions.
- A unique event is now generated that notifies when the RBA Engine has a backlog of User-Initiated tasks and indicates when the engine is unable to pull new tasks for execution due to in-progress execution.

- Filter-while-you-type fields are now available for the "Aligned" fields within the Run Book Automation Policy Editor.
- The Automation Engine can now report in-memory queue sizes every minute.
- The Automation Engine now creates an internal alert ("793: The automation engine is still processing %V") instead of a critical system log when the Automation Engine exceeds thresholds.
- The "BETA" label was removed from the Action Type Manager page (Registry > Run Book > Action Types) title and menu navigation option.
- The data retention threshold was synced to interact correctly with the appropriate Run Book Automation queue-size metrics.

Schedules

- Schedule data from the SL1 Central Database host will now be pruned based on the "Schedules Purge Timeout" setting. After the time expires, the schedule will be removed from the database.
- Users can now select the "Preserve Schedule" checkbox in the Schedule Editor of the SL1 classic user interface to exclude specific schedules from being pruned after expiration.

ScienceLogic Support Pack

- The "Support: Concurrent PowerShell Enabled Config" and "Support: Normalized Tables Size Performance" Dynamic Applications were added to the ScienceLogic Support Pack PowerPack. The former collects information about stacks that are using the PowerShell collector service, and the latter collects information about the size of database normalized tables.
- The "Support: SL1 Performance" Dynamic Application was updated to provide memory statistics to determine memory congestion and fragmentation on an SL1 system.
- The "Support: Service Usage Config" Dynamic Application was updated to collect availability from other services.
- VMware self-monitoring was added to the ScienceLogic Support PowerPack.
- The "Support: MySQL Performance" Dynamic Application was replaced with the "Support: MariaDB Performance" Dynamic Application.

Security

- On newly installed systems of 11.2.0 or greater, SL1 now defaults to HTTPS. You can change this setting in the "Force Secure HTTPS" field on the Behavior Settings page (System > Settings > Behavior).
- 11.2.0 includes package updates that improve security and performance.
- Updated packages to address two "high" security vulnerabilities, ELSA-2022-9227 and ELSA-2022-9421.

SL1 Appliances

- A new Nodes page (Manage > Nodes) displays key properties and metrics for all nodes, including SL1 collectors, SL1 instances, and SL1 Database Servers. This new page includes an "Add Node" button at the top right of the page that lets you connect new SL1 Data Collectors, both Traditional and PhoneHome Collectors. For more information, see the "Installing and Configuring an SL1 Collector" chapter in the [Installation](#) manual.
- On this new Nodes page, when you click the "Add Nodes" button to connect a new Data Collector with the SL1 Database Server, a three-step wizard displays that lets you choose the connection type, define the collector properties (including the collector name, IP address, SL1 Collector Group for the new collector, and collector type), and then generate a token to establish the connection between the collector and the database. This wizard replaces the steps that require you to use the "Web Configuration Utility" in previous versions of SL1. The PhoneHome page in the Web Configuration Utility contains a link that redirects you to the Nodes page in SL1. For more information, see the "Installing and Configuring an SL1 Collector" chapter in the [Installation](#) manual.
- On the Nodes page, you can click the name of an existing node to open a modal that enables you to edit details for that node.
- The Nodes page also includes a Pending tab for tokens that you can use to establish a connection between a collector and an SL1 Database Server, and a Tokens tab that displays a list of existing and expired tokens used for connecting nodes. On the Tokens tab, you can recreate an expired token, which lets you generate a new token with the same collector details. Recreating the token actually deletes the existing token but uses its user-supplied collector registration details to generate a new token.

SL1 Extended Architecture

- SL1 Extended Architecture now includes the Prometheus AlertManager. This program supports internal alerting.

Snippet Framework

- A REST Toolkit enhancement was added to the Snippet Framework that adds logging capability to silo-authentication. This allows any Token Authentication users to debug any request/credential issues, which can be used to troubleshoot issues during Dynamic Application development.
- A new command-line enhancement was added to the Snippet Framework that allows collection objects to be retrieved with either an agent, by running the command directly on the host in a local shell, or remotely over SSH.

System Status Script

- Made several updates to the `system_status.sh` script to fix central database and data engine detection logic, adjust the logic to check for broken symlinks on collector appliances, and modify license check and license status output logs for Message Collectors and Data Collectors.

System Update

- Deployment log messages now provide additional details about why a pre-upgrade failure might be occurring.

Telemetry

- Metrics and usage data were added to the Collector Pipeline service to better troubleshoot problems and more accurately collect billing data.
- Metrics were added to the Enterprise Database: Collector Config Push process (config_push.py). This process "pushes" information from the SL1 Database Server to the Data Collectors and Message Collectors.

User Interface

- The following inventory pages and tabs in the SL1 user interface were updated to use a new set of filters for the columns in the list. You can start typing filter text or select filter options in one or more of these filters to narrow down the list to just the items you want to view:
 - Agents (Devices > Agent)
 - Business Services
 - Collector Groups (Manage > Collector Groups)
 - Dashboards
 - Devices
 - Device Investigator
 - Changes
 - Events
 - Interfaces
 - Investigator
 - Logs
 - Ports
 - Processes
 - Software
 - Device Categories
 - Device Classes
 - Events
 - Event Policies
 - Event Message

- Suppression
 - Global Manager Stacks
 - Nodes
 - Subscription Usage
 - License Usage
 - License Config
- On the new inventory pages, column attributes such as width, order, sort, and refresh interval will persist after you adjust them. You can click **Reset** to return the columns to their default display settings.
 - From the **Grid Settings** menu on the new inventory pages, you can click **Clear Filters** to remove any column filters on that page.
 - Some of the inventory pages now list the number of items on the page at the top of the page.
 - If you type invalid search terms or incorrect syntax in a filter on an updated inventory page, SL1 will highlight the headers in red. SL1 will ignore the incorrect input until you fix it.
 - Clicking the **Help** button now opens the **Help** menu on the right-hand side of the SL1 user interface. The **Help** menu also includes a context-sensitive link to the relevant page in the product documentation.
 - The labels in the search icons (criteria buttons) were updated to be more readable. For example: "property" instead of "className.property".
 - At the top of the **Events** page, the severity filters now display as oval-shaped "pills" instead of boxes.

VMware Tools

- If the `libtool-ltdl` and `libxslt` packages were already installed on the host where you run the `yum downloadonly` command to install VMware Tools, those dependent packages will be listed as missing and will prevent you from installing VMware Tools. To work around this, remove those two packages from the host before installing VMware Tools.

Webhooks

- You can now use webhooks as a method for streaming alerts and messages from a device to SL1. With this feature, an external system (the "provider") sends webhook messages via HTTP to be ingested by an SL1 Message Collector (the "receiver"). To use this feature, you must add an addressable webhook receiver and then align that receiver to the device from which you want to ingest webhook messages. You can add, edit, or delete webhook receivers on the new **Webhook** page (**Registry > Monitors > Webhooks**) in the classic SL1 user interface. For more information, see the section on "Using Webhooks to Generate Events" in the [Events](#) manual.

WMI

- SL1 collection via WMI Dynamic Applications was refactored in response to Microsoft security updates. WMI Dynamic Applications can now collect data from hardened Windows servers.

CAUTION: This change significantly decreases the number of Microsoft Windows servers that can be supported on each Data Collector in your SL1 system. Users who need to monitor Windows devices using WMI should analyze their system resources and capacity before upgrading to 11.2.0. For guidance about sizing, see the updated [Collector Sizing guidelines for WMI endpoints](#).

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

Issues Addressed in 11.2.0

Account Information

- Addressed an issue where the Time Zone value was changed when a user updated information on the Account Information page (Preferences > Account > Information). (Case: 00195148) (Jira ID: EM-46184) (Jira ID: EM-46184)

Agent

- Addressed an issue where Agent collection could not be turned off during device maintenance windows. (Case: 00137491) (Jira ID: EM-41848) (Jira ID: EM-41848)
- Addressed an issue that was causing various fields to not update on SL1 Message Collectors whenever an Agent was uninstalled from a device. (Case: 00172800) (Jira ID: EM-44078) (Jira ID: EM-44180) (Jira ID: EM-44180)

Assets

- The Disk Array Size field on the Asset Configuration page (Registry > Assets > Asset Manager > create button/wrench icon > Configuration tab) is no longer truncated. (Case: 001908) (Jira ID: EM-445683) (Jira ID: EM-45646) (Jira ID: EM-45683)

Backups

- Addressed an issue that was causing the backup retention database tool command to work only on full backups and not on configuration backups. (Case: 00213052) (Jira ID: EM-47564) (Jira ID: EM-47580) (Jira ID: EM-47580)
- Addressed an issue that was causing configuration backup to periodically fail when trying to back up the temporary table "master_dev.topo_comp_depth_rebuild". With this update, that particular temporary table was added to the list of tables that are excluded in a configuration backup. (Case: 00188441) (Jira ID: EM-45371) (Jira ID: EM-45371)

Business Services

- Addressed an issue where, during the creation of a Business Service, one of the metrics required users to pick an Index despite being non-Indexed in the classic SL1 user interface. This resulted in failed aggregation. (Case: 00160391) (Jira ID: EM-43364) (Jira ID: EM-43364)
- Addressed an issue in which any device search that included the string "North America" took more than two minutes to return results on the Business Services page. (Case 00213380) (Case: 00232243) (Jira ID: EM-47549) (Jira ID: EM-47435) (Jira ID: EM-47549)
- Addressed an issue with slow GQL queries that affected Business Services. Queries for devices based on custom attributes no longer run slowly, hang, or cause outages. (Case: 00232092) (Jira ID: EM-49043) (Jira ID: EM-49045) (Jira ID: EM-49216) (Jira ID: EM-49224) (Jira ID: EM-49225) (Jira ID: EM-49224)

Collector Groups

- Addressed an issue in which devices that were put into virtual collector groups (VCUGs) were causing deviation crunch policies to spawn an asynchronous job for every collector, despite VCUGs not being enabled for collection. With this update, SL1 no longer spawns any asynchronous da_hourly_stats_update tasks for deviation crunch policies associated with VCUGs. (Case: 00170293) (Jira ID: EM-44118) (Jira ID: EM-44170) (Jira ID: EM-44171) (Jira ID: EM-44171)

Credentials

- Addressed an issue that was enabling users to delete credentials that were already assigned to devices, which then caused those devices to disappear from the Device Registry page (Registry > Devices > Device Manager). With this update, SL1 will now display an error message if you attempt to delete a credential that is assigned to devices. Also improved how PowerPacks and Device Templates check and use credentials. (Case: 00189403) (Case: 00189403) (Jira ID: EM-46085) (Jira ID: EM-47189) (Jira ID: EM-45719) (Jira ID: EM-46085)
- Addressed an issue in which PowerShell credentials were not translating the variable %N as the device hostname, as intended. (Case: 00147853) (Jira ID: EM-41741) (Jira ID: EM-41817) (Jira ID: EM-41817)
- Addressed an issue that was enabling users to delete credentials that were already assigned to devices, which then caused those devices to disappear from the Device Registry. With this update, SL1 will now display an error message if you attempt to delete a credential that is assigned to devices. (Case: 00189403) (Jira ID: EM-45719) (Jira ID: EM-47189)(Jira ID: EM-47189)

Dashboards

- An enhancement was introduced in which the Service Type drop-down menu within a Service widget is now hidden when the widget consumes context. (Case: 00084743)(Jira ID: EM-36018)
- After installing the "F5 BIG-IP: SL1 Dashboards" PowerPack on an SL1 system that has been built or upgraded with an ISO, the dashboards now appear on the Dashboards page. (Case: 00139163) (Jira ID: EM-42101, EM-40951)
- Optimized a dashboard-related query that was affecting user interface performance. (Case: 00204331) (Jira ID: EM-46852)

- Addressed an issue in which classic user interface Leaderboard widgets failed to load and configure data properly after the PHP 7 upgrade in SL1 11.1.0. (Case: 00219319) (Jira ID: EM-48628)

Device Groups

- The "Software Title" and "Software Version" rule selectors in the Device Group Rule Editor have been renamed "Asset Software Title" and "Asset Software Version", respectively, to more precisely identify what collected properties of devices are used to match on rules using these selectors. (Case: 00190190) (Jira ID: EM-46292, EM-45309)

Device Templates

- Device templates with multiple log policies can now be applied to one or more devices. (Case: 00158562) (Jira ID: EM-42753)
- Applying a new device template with a new organization name to a device will now update the organization associated with active SL1 events related to that device. (Case: 00168636) (Jira ID: EM-44128)

Device Tools

- Updated an incorrect command in the Device ARP Ping tool to ensure the tool works properly. (Case: 00039376) (Jira ID: EM-45072)

Devices

- When a device is manually deleted, vanished, or purged, extended-type custom attributes will be removed from the SL1 central database. (Case: 00226938) (Case: 00242287) (Jira ID: EM-49869) (Jira ID: 50073) (Jira ID: EM-50127)
- An issue was addressed in which deleting a device would not delete its app_meta data in Dynamic Applications. (Case: 00195306) (Jira ID: EM-45407)
- SL1 now reviews the file systems detected on devices and updates the type if necessary. It will also purge previously collected file systems from the hardware table if they are not collected within 24 hours. You can change this default 24-hour purge interval by updating the device_hardware_purge_hours value on the master.system_custom_config table. (Case: 00170155) (Case: 00210669) (Jira ID: EM-44620) (Jira ID: EM-44669)
- When using a Windows Agent, updated device hostnames will now be reflected in SL1. (Case: 00207059) (Jira ID: EM-46973) (Jira ID: EM-46994)

Discovery

- In the Discovery Session Editor page (System > Classic Discovery > create button or wrench icon), the currently saved values in the SNMP Credentials, Other Credentials, and Collection Server PID fields are now surrounded by brackets. (Case ID: 00216440) (Jira ID: EM-47745) (Jira ID: EM-47772) (Jira ID: EM-47772)

- Addressed an issue that was causing the discovery process to enter a loop when it encountered a device class identifier with no match in the system. With this update, SL1 will now assign a device to the Generic SNMP device class if it cannot match the device to another device class. (Case: 00184055) (Jira ID: EM-45153) (Jira ID: EM-45237)(Jira ID: EM-45237)
- Addressed an issue that was causing discovery to fail when users attempted to run, terminate, and then re-run a discovery session shortly after upgrading SL1. With this update, the discovery will run successfully. (Case: 00182836) (Jira ID: EM-44801) (Jira ID: EM-44869) (Jira ID: EM-44869)
- Addressed an issue that was causing device classes to be incorrectly defined during discovery when the discovery identifier (SysObjectID) included a leading period. With this update, device classes that include a leading period are now discovered correctly. (Case: 00205198) (Case: 00216454)(Jira ID: EM-46490)

Documentation

- Updated the [Monitoring with the SL1 Agent](#) manual to specify that the ScienceLogic Agent for Linux now adds a new user "scilog" and a new group "scilog" to the operating system. (Case: 00225286) (Jira ID: EM-48510)
- Updated the [Run Book Automation](#) manual with the complete list of credential types, and to address an issue with the Snippet Action Dynamic Application credential example documentation. (Case: 00236324) (Case: 00245395) (Jira ID: EM-49817) (Jira ID: EM-49737) (Jira ID: EM-50452) (Jira ID: EM-50725)
- Updated the [Run Book Automation](#) manual to add missing xtype and ytype info in the section "Writing Snippet Code". (Case: 00217838) (Jira ID: EM-47908)
- Updated the [ScienceLogic Support Pack PowerPack](#) manual to note that the "Support: MySQL Performance" Dynamic Application has been replaced with the "Support: MariaDB Performance" Dynamic Application. (Jira ID: EM-49766)
- Updated the instructions in the [System Administration](#) manual about changing root passwords on Database Servers and Data Collectors to reflect changes made in MariaDB 10.4. (Case: 00116783) (Jira ID: EM-39167)
- Updated the [System Administration](#) manual with additional information about the features and fields on the Appliance Manager page (System > Settings > Appliances). (Case: 00231891) (Jira ID: EM-49034)

Dynamic Applications

- A query executed by the Dynamic Application "Mail Backlog" in the ScienceLogic Support PowerPack no longer generates false alerts. (Case: 00208019) (Jira ID: EM-47110) (Jira ID: EM-47285)
- Addressed an issue in which some non-snippet Dynamic Applications that were not aligned with execution environments were generating system log errors during regular or nightly discovery. With this update, these errors will no longer appear for Dynamic Applications that are not aligned with execution environments or assigned to PowerPacks. (Case: 00111906) (Case: 00112429) (Case: 00112718) (Case: 00115486) (Case: 00118593) (Case: 00121302) (Case: 00121570) (Case: 00120298) (Jira ID: EM-38807) (Jira ID: EM-38808)
- In the Threshold Objects page (System > Applications > create button/wrench icon > Thresholds tab), the Numeric Range High field and the Numeric Range Low field now allow up to 13 digits, with 3 digits after the decimal point. (Case: 00192248) (Jira ID: EM-46032) (Jira ID: EM-46042)

- If a device uses multiple Dynamic Applications and those Dynamic Applications all assign a value to a single collection label, and all but one of the Dynamic Applications are unaligned from the device, the remaining Dynamic Application will continue to populate the collection label. If a Dynamic Application that assigns a value to a collection label is unaligned from a device, the Collection Labels page (System > Manage > Collection Labels) will immediately update the Aligned Devices field; previously, this could take up to one hour. (Case: 00093389) (Jira ID: EM-36891)
- Addressed an issue that was causing Dynamic Applications that were using the Index type collection objects to generate a "Non-digits in oid" error. (Case: 00190863) (Jira ID: EM-45315)
- A database view was optimized and indexed to prevent 504 errors when loading the list of subscribers for a Dynamic Application in the SL1 user interface and via the REST API. (Case: 00178865) (Jira ID: EM-44939) (Jira ID: EM-44886) (Jira ID: EM-44924) (Jira ID: EM-44929)
- Updated the subscribers count query to improve the performance of the Dynamic Applications page (System > Manage > Dynamic Applications). (Case: 00235083) (Jira ID: EM-49263)

Email

- Ensured that images in inline HTML reports in email display properly. (Case: 41830) (Jira ID: EM-41706) (Jira ID: EM-41830)

Email Round-Trip Monitoring Policies

- Addressed an issue where Entity Name did not update within an Email Policy when the name of the aligned default entity was changed after being assigned to the Email Policy. (Case: 00209991) (Jira ID: EM-47584)
- Addressed an issue where Organization ID was not updated within an Email Policy when the Organization was changed for the default entity assigned to the Email Policy. (Support ID: 142243) (Jira ID: EM-21838)

Events

- Addressed an issue in which some healthy and unhealthy events were processed incorrectly due to incorrect specified order. (Case: 00206646) (Jira ID: EM-47041) (Jira ID: EM-47546)
- When an event policy is deleted from SL1, the event is also removed from all Run Book Automations. (Case: 0192943) (Jira ID: EM-45669) (Jira ID: EM-45747) (Jira ID: EM-14377)
- Addressed an issue that was causing the "EM7: DRBD Failover Primary" event policy to falsely trigger an alert when the relevant collection objects registered as "0" (zero) for a polling cycle due to a missed collection or failover. With this update, the DRBD Failover Primary alert formula now checks if the prior value is non-zero to ensure that it does not alert if no data exists. (Case: 00125462) (Case: 00128562) (Case: 00130154) (Case: 00132711) (Case: 00144544) (Jira ID: EM-40121) (Jira ID: EM-40168)
- In events about IT Services (in the classic SL1 user interface), if the user account that created the IT service has been deleted from SL1 or is missing, SL1 provides a default value for xname and roa_id. (Case: 00037424) (Jira ID: EM-31148) (Jira ID: EM-32250)

File Systems

- Addressed an issue with ICDA and file system inventory and performance collection. File systems are now discovered and monitored. (Case: 00185984) (Jira ID: EM-45537) (Jira ID: EM-46173)

GraphQL

- Addressed an issue where GraphQL responses contained extremely large cursor sizes. A cursor contains the metadata for the ID of an element, and the cursors in GraphQL were taking up an unexpectedly large amount of memory. (Case ID: 00240115) (Jira ID: EM-50193)
- The GQL query for deviceRelationships now returns the correct DCM+R relationship type. (Case: 00227848) (Jira ID: EM-49546) (Jira ID: EM-459503)(Jira ID: EM-50353)(Jira ID: SLUI-14078)

High Availability and Disaster Recovery

- In certain failure scenarios for disaster recovery (DR) backups, the suppressed DRBD events will not be unsuppressed. The system will now periodically check for these orphaned suppressions and unsuppress the events. (Support ID: 90231) (Jira ID: EM-7844)
- Addressed an issue that was preventing successful failover between primary and secondary servers when MySQL took longer than 20 minutes to shut down, due to a default Pacemaker process setting that overrode the MySQL shutdown timer after 20 minutes. With this update, the Pacemaker shutdown-escalation time value was adjusted to match the MySQL timeout value. (Case: 00132827) (Jira ID: EM-40307)

Internal Collections Dynamic Applications

- An issue was addressed in which Internal Collections Dynamic Applications (ICDAs) were not working for large volumes. If a volume is larger than 16TB, ScienceLogic recommends that you use operating system-based PowerPacks instead because they can reliably gather the information for such volumes. (Case: 00211635) (Jira ID: EM-47497) (Jira ID: EM-47612)
- Addressed an issue that was causing data collectors to terminate collection without completing (SIGTERM) when attempting to load balance component devices that were aligned to ICDAs with default collector affinity. With this update, an ICDA with default collector affinity that is aligned to a component device no longer forces the device to be assigned to its root device collector. (Case: 00126258) (Jira ID: SOL-16498) (Jira ID: EM-41250)

Licensing and Billing

- Updated the default retention settings (System>Settings>Data Retention) for Subscription Usage Data. (Case: 00221390) (Jira ID: EM-48615) (Jira ID: EM-48089)

Logging

- Addressed an issue that was causing /var/log to exceed the recommended monitoring threshold on larger environments. With this update, /var/log has been resized for larger virtual machines. (Case: 00187197) (Jira ID: EM-45071) (Jira ID: EM-45159)

Login Page

- On the login page, when a user clicks the "I forgot my password" link, SL1 now either sends an email to the user with the link to reset the password, using the IP address of the Administration Portal instead of "localhost"; or displays the instructions to reset the password, for users without an email address in SL1. (Case: 00218846) (Case: 00180421) (Jira ID: EM-48017) (Jira ID: EM-45272) (Jira ID: EM-44709) (Jira ID: EM-47203)

Log Monitoring Policies

- An issue was addressed in which backslashes in log monitoring policy filters resulted in log monitors that did not monitor files properly. (Case: 00151985) (Jira ID: EM-41805)

Maps

- Addressed an issue in which Map Nodes failed to query when the map's seed query was left empty. Additionally, Map Images failed to add as well. Now, the seed query, even if left empty, will return nodes. (Case: 00190138) (Jira ID: EM-45346)

Network Interfaces

- Addressed an issue in which new Interface types that were created using the Interface Type Editor page (System > Customize > Interfaces) had an IANA Identification Number of "0" by default. With this update, new interface records will have auto-incrementing IANA Identification Numbers. (Case: 00202930) (Jira ID: EM-46418) (Jira ID: EM-47148)
- Addressed an issue in which interface rediscovery was causing interfaces without MAC addresses to be swapped with other, random interfaces in Bandwidth Billing policies. With this update, new records will now be created for interfaces that are discovered without MAC addresses. (Case: 00168030) (Jira ID: EM-43928) (Jira ID: EM-44378)

PhoneHome Collectors

- Addressed an issue that was causing a mismatch between the Data Collector's PhoneHome user's known_hosts file and the Database Server's public authentication key after rebuilding a Database Server, which in turn was causing a connection failure to the Database Server. (Support ID: 103182) (Jira ID: EM-12247)
- Addressed an issue that was causing a new SSH key to be generated whenever an Azure SL1 appliance was rebooted, which in turn required PhoneHome collection to be reconfigured. (Support ID: 174310) (Case: 00015425) (Jira ID: EM-28109)

Platform

- To address an issue that was causing swap utilization to continually grow after upgrading SL1, Transparent Huge Pages (THP) are now disabled on systems that have Extensible Firmware Interface (EFI) enabled. (Case: 00203966) (Jira ID: EM-46521) (Jira ID: EM-47609)

- After enabling Force Secure HTTPS in the Behavior Settings page (System > Settings > Behavior), the header for "Strict-Transport-Security" no longer includes quotation marks that make the header invalid. (Case: 00178812) (Jira ID: EM-44923)
- Data pruning for Data Collectors was failing due to missing values in the table in `_storage.data_storage_last_collect`. Removing or truncating values from that table is not a part of normal operation. A log message with Minor severity now shows up in the System Log if data is missing from this table. (Case: 00045117) (Jira ID: EM-31853)
- Addressed an issue in which the outbound messaging process failed when certain invalid characters were present. (Case: 00113195)(Jira ID: EM-28879)
- Addressed an issue where the `apl-aggregator-vitals` service was restarting and eventually failing when it could not find a system file. (Jira ID: AP-1827)
- The default memory limit that the Daily Maintenance process in SL1 can use was increased from 1GB to 2GB to allow processing of more rows to query and delete. This should improve the pace of pruning of old data. (Case: 00118596) (Jira ID: EM-39482)
- Addressed an issue where the Vitals Aggregator topic grew too large in an environment with a large number of services or processes, which prevented the SL1 system from scaling. (Case: 00172382) (Jira ID: EM-44127)
- Removed nullok from system configuration files governing authentication, and password complexity for the local user account. (Jira ID: EM-41978)
- The default values for "pruner_MinRowLimit" and "pruner_initialRowLimit" were updated from 10000 to 3000 as there are now two pruners running at the same time to alleviate the buildup of logs. (Case: 00213388) (Jira ID: EM-47438)
- An issue was addressed in which asset sync was failing due to 128k entries appearing in the `master_dev.device_hardware` table for a single device. To address this, there is now a new Daily Maintenance task called "purge master_dev.device_hardware stale rows", which is disabled by default. To enable this process, update the "is_enabled_stale_hard_maint_task" setting to true (1). (Case: 00100023) (Case: 00161829) (Case: 00136335) (Jira ID: EM-37732)
- Improved SQL query in the data access layer to provide the Event `xname` and `roa_id` if IT Service Owner is missing. (Support ID: 142020) (Jira ID: EM-46853)
- Addressed an issue that was causing the "EM7 Core: Python Envs Clean" system process to occasionally experience unhandled exceptions. This was occurring when the process encountered a broken symlink or a race condition while looking for expired environments to be removed. With this update, instead of raising the exception, the process logs a message if it encounters a broken symlink and then continues, or tries again after it has finished processing the other files in the directory if it encounters a race condition. (Case: 00175062) (Jira ID: EM-45489) (Jira ID: EM-45579)
- Addressed an issue in which the `master.policies_device_monitors` table was not matching the potential maximum SL1 device ID value. (Case: 00194765) (Jira ID: EM-45084) (Jira ID: EM-45222)
- Addressed an issue that was causing empty `php-error` logs to generate daily. These logs then failed to rotate, which in turn caused the files to grow so large that they were causing file system usage warnings. With this update, both the logging and log rotation behaviors were addressed, and logs are now rotating regularly. (Case: 00089934) (Case: 00101351) (Case: 00159318) (Case: 00181015) (Jira ID: EM-37039) (Jira ID: EM-43401) (Jira ID: EM-44700) (Jira ID: EM-44713)

PowerPacks

- Addressed an issue that was causing the task manager to repeatedly crash when attempting to use the Cisco: Wireless v102 PowerPack to discover Cisco WLCs with a large number of access points, due to a large backlog of asynchronous jobs. This update improves the performance of long-running queries that insert asynchronous Dynamic Application configuration records into the database server by bundling these records together and inserting them into one query where possible rather than having a query for each individual record. The queries are executed in the 'pre_proc_config' module's 'push_dynamic_app_config()' function. (Case: 00189110) (Jira ID: EM-45300) (Jira ID: EM-45839)
- Addressed an issue where missing records in master_filestore.file_index for ScienceLogic Libraries prevented a PowerPack installation. (Case: 00085430) (Jira ID: EM-35823) (Jira ID EM-39250)
- Made improvements to the performance of legacy PowerShell Dynamic Applications. If a PowerShell command takes longer than 90 seconds to execute, it will be terminated and the collection data for that execution will be lost. This will prevent collection from hanging and allow other PowerShell commands to be executed. Only the individual command that takes longer than 90 seconds will be terminated; this does not affect any other commands from the same Dynamic Application or from other Dynamic Applications. In addition, if a connection to a Windows server is closed without closing the PowerShell command, the orphaned PowerShell command will continue to exist on the Windows device for two minutes and then will automatically close. (Case: 00100829) (Jira ID: EM-37463) (Jira ID: EM-47835) (Jira ID: EM-36579)
- The "SL1: Concurrent PowerShell Monitoring" PowerPack was created to parse and classify Concurrent PowerShell Logs. This PowerPack contains Dynamic Applications that parse logs and trigger alerts according to severity. (Case: 00154016) (Jira ID: EM-42313)

REST API

- The ScienceLogic API now allows the device/device_ID/class_type resource to support either a GUID or a device class ID. (Jira ID: EM-45367)

Run Book Automations

- Addressed an issue where the right single quote character, when utilized in an outbound email message from Run Book Automation, caused an unhandled exception. (Case: 00113892) (Jira ID: EM-39031) (Jira ID: EM-39082)
- Non-existent event policies in the Run Book Automation policy's event table were removed to ensure that the Automation Policy Manager page (Registry > Run Book > Automation) displays the correct number of linked event policies. (Case ID: 00192943) (Jira ID: EM-45669) (Jira ID: EM-45748)
- Addressed an issue with "Send an email notification" Run Book Action types. If an email address is aligned to a "Send an Email Notification" Run Book Action type and that email address is removed from SL1, the email address is no longer included in the Run Book Action. (Case: 00098608) (Jira ID: EM-37960)

Schedules

- An issue was addressed in which schedules in the classic user interface were not working for monthly iterations. (Cases: 00103698 and 00163277) (Jira ID: EM-39499)
- An issue was addressed in which schedules with start times in GMT +/- N time zones would appear incorrectly in the user interface. (Case: 00121291) (Jira ID: EM-40046)

ScienceLogic Support Pack

- In the ScienceLogic Support PowerPack, updated the "Support: MariaDB Performance" Dynamic Application to generate an event when the table open cache or table definition cache are undersized for the current workload. (Case: 00201293) (Jira ID: EM-46144)

SNMP Walk

- Addressed an issue that was causing SNMP walks to fail due to the inclusion of single quotes (') or double quotes ("). (Case: 00158636) (Jira ID: EM-43439)

System Status Script

- Resolved an issue with the system_status.sh check where Message Collectors are shown as unlicensed when they are in fact self-licensed after integration into a distributed SL1 stack. (Case: 00206422)(Jira ID: EM-46742)
- Addressed an issue that was causing the silo_mysql and system_status.sh scripts fail to connect to the database when some special characters were included in the dbpasswd field in the silo.conf file. This update improves the handling of special characters for database access. (Case: 00203158) (Jira ID: EM-46398) (Jira ID: EM-46813)
- Addressed an issue that occurred during detection routines for local and remote databases. (Case: 00217969) (Jira ID: EM-47787)
- Addressed an issue with the system_status.sh script. The script no longer reports erroneous broken symlinks for Data Collectors. (Case: 00217965) (Jira ID: EM-47912)

System Update

- Addressed an issue where upgrades caused the error "error checksumming master.system_credentials" and "error checksumming master.system_credentials_snmp" to appear in the System Logs page (System > Monitor > System Logs). (Case: 00219536) (Jira ID: EM-48172) (Jira ID: EM-47985)
- Addressed an issue that was preventing some containerized service processes from stopping during upgrades. With this update, any leftover processes that should not be running during a Docker update will be stopped, and SL1 will log diagnostic information if failures are encountered during upgrading. (Case: 00223991) (Case: 00239398) (Jira ID: EM-48321) (Jira ID: EM-48340) (Jira ID: EM-49161)

User Interface

- Fixed a CSS error in the Date Range Selection fields throughout SL1. Fields are now properly aligned. (Case: 00208856) (Jira ID: EM-47132) (Jira ID: EM-47113)
- Addressed an issue where filtering by Device Categories or Device Classes in the Device Browser was not working. (Case: 00012022) (Jira ID: EM-27890)
- Addressed an issue where searching device classes on devices sometimes showed invalid results. (Case: 00233392) (Jira ID: EM-49179)

Web Content Policies

- Addressed an issue with Web Content policies. Web Content policies with URI that cannot be resolved no longer generate "success" entries in device logs. (Case: 00112098) (Jira ID: EM-39213) (Jira ID: EM-39109)

Windows Service Policies

- Addressed an issue where a PHP/MySQL error occurs when you try to add a new Windows Service Policy. (Case: 00217578) (Jira ID: EM-48880)

SL1 Extended Architecture

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

NOTE: 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.

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

PowerPacks in 11.2.4.1

Before upgrading to 11.2.4.1, 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 PowerPacks you do not want to update and select their checkboxes.
3. Click the **Select Action** field and choose *Change Collection State: Disabled (recursive)*, and then click the **[Go]** button.
4. Wait five minutes after disabling collection.
5. Install the SL1 update.
6. Go to the **Device Components** page (Registry > Devices > Device Components).
7. Select the checkbox for all affected root devices.
8. Click the **Select Action** field and choose *Change Collection State: Enabled (recursive)*, and then click the **[Go]** button.

New and Updated PowerPacks

The 11.2.0 release **includes** the following PowerPacks that are new or updated:

- Microsoft: Windows Server version 113

PowerPacks Removed from ISO Due to New Pricing

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

NOTE: If you are upgrading from a previous version of SL1, the 11.2.4.1 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.

Upgrade Process for Systems Running 8.1.0 and Earlier

WARNING: ScienceLogic strongly suggests you contact Customer Support or your Customer Success Manager to plan your migration from CentOS (versions of SL1 prior to 8.1.1) to 11.2.4.1.

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 firewalld **before you upgrade**. Please contact ScienceLogic Support for more information.

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

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

To download updates for previous SL1 software versions that have reached their End-of-Life date and are no longer supported by ScienceLogic, contact ScienceLogic Support or a designated Customer Success Manager to get the update files.

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

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

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

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

Downloading SL1 Updates on SL1 Systems Running 8.6.0 or Later

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

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

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

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

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

Recently Deprecated Features

11.2.0

- The Inbox feature has been deprecated and is available only for reports in the classic user interface.
- The IPMI features have been deprecated in SL1. To re-enable this feature, contact ScienceLogic Customer Support for instructions.
- The Knowledge Base and Knowledge Base tab, which have been disabled by default since SL1 8.10, have been removed from SL1 11.2.0.
- Related to the deprecated Knowledge Base, the Clipboard feature was removed from Ticketing. The access hook for the clipboard feature were also removed.

11.1.0

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

10.1.0

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

© 2003 - 2024, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

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

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

Copyrights and Trademarks

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

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

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

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

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

Other

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

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

ScienceLogic

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

International: +1-703-354-1010