Skylar One 12.5.20 Release Notes

Skylar One version 12.5.20

Skylar One 12.5.20 Release Notes

IMPORTANT: ScienceLogic strongly recommends that you review the *installation and upgrade instructions*, *important upgrade notes*, and *known issues* for this release before installing or upgrading to Skylar One 12.5.20.

The Skylar One 12.5.20 release includes the following new features and enhancements:

- *New versions of the Linux and Windows agents*
- *Enhanced topology settings for devices*
- *A redesigned Event Suppressions page*
- *Geographic maps in Business Services, plus a new Locations page*
- *A new PowerPack security and signing framework*
- *Improvements to the Anomaly Detection page*
- *Significant improvements to the system upgrade process*
- *Support for ARP-based topology relationships and numerous other topology enhancements*
- Plus *many additional updates* and *addressed issues* from previous releases

These release notes provide a comprehensive list of the features, enhancements, and addressed issues that are included in the Skylar One 12.5.20 release.

NOTE: Some of the features, enhancements, and addressed issues that are included in this Skylar One platform release were originally included in the following Skylar One AP2 releases:

- *8.29.41 (Nougat)*
- *8.34.20 (Opera Cake)*
- *8.34.116 (Pancake)*

AP2 version 8.34.116 (Pancake) is installed by default in Skylar One 12.5.20.

This document covers the following topics:

Before You Proceed

If you are planning to consume Skylar One 12.5.20, be advised of the following:

- The 12.5.20 release is available as both an ISO and patch.
- All Python 2 functionality was removed from Skylar One with the 12.5.1 release. If you still use Python 2 for custom code, you cannot upgrade to this or any later releases until your custom code is Python 3-compatible.
- You can perform upgrades from one minor version to any later minor version within the 12.5.x series. ***However, as with all updates, ScienceLogic strongly recommends that you perform such upgrades in a test environment before implementing the upgrades in production environments.***
- The ScienceLogic-validated upgrade paths to version 12.5.20 are outlined below.
    - 12.5.7 > 12.5.20
    - 12.5.5 > 12.5.20
    - 12.5.4 > 12.5.20
    - 12.3.14 > 12.5.20
- You cannot upgrade to 12.5.20 from version 12.5.8 due to release timing.
- 12.5.20 supports STIG deployments. However, STIG-compliant users should contact their ScienceLogic account managers for important information about upgrading to this release.
- AWS deployments that are using Aurora 3 can upgrade to this release.
- The Enterprise Key Management Service (EKMS) is enabled by default in 12.5.20.
- There is a known issue impacting upgrades to 12.5.4 and later from the 12.3.4 through 12.3.6 releases that could result in deployment failure due to patch hook task issues. For more information about this issue as well as a workaround, see https://support.sciencelogic.com/s/article/20990.
- When upgrading Skylar One to version 12.5.20 from 12.3.x or earlier, PowerPacks that use legacy encryption will become read-only. This issue does not impact upgrades from previous 12.5.x releases. For more information about this issue, including resolution steps, see https://support.sciencelogic.com/s/article/20806.(Jira ID: EM-73018)
- If you currently use local authentication with an MD5-hashed password, you will be required to reset your password when upgrading to Skylar One 12.5.20.
- 12.5.20 adds a new, post-update cleanup script that removes old patch bundles to free space on the Database Server. This cleanup process can increase upgrade time during the first run, especially on systems with a large number of older patch bundles. To reduce upgrade time, you can pre-clean older patches before upgrading. For more information, see the section on System Update Enhancements.

For more information, see the *Important Upgrade Notes* and *Known Issues* sections.

New Features and Enhancements in Skylar One 12.5.20

This section describes the new features and enhancements that are included in Skylar One 12.5.20.

Agent

- **What's new:** ***New versions of the Linux and Windows agents.*** Skylar One version 12.5.20 adds Linux Agent v198 and Windows Agent v155, which include the following enhancements since the previous versions that were pinned to a Skylar One release in version 12.5.1:
    - Linux Agent v197:
        - Improved the handling of polled data processes to prevent defunct processes created by polled data commands. (Case: 00504475)
    - Linux Agent v198:
        - Updated AIX dynamic linking so the same agent executable runs on both AIX 7.2 and 7.3. (Case: 00584663)
        - Added transaction summaries for non-intercepted processes and transactions.
        - Specified a `nologin` shell when creating the `scilog` user during installation.
        - Updated the agent to use the `MemAvailable` field from `/proc/meminfo` to report RAM usage, when available.
        - Improved the agent log timestamp format.
        - Updated the default value of the `DataKeep` parameter from 1 to 10.
        - When file systems change size or RAM is added, it now triggers the generation and uploading of a new system file.
        - Added a sync every three hours for Dynamic Applications aligned to the agent devices.
        - Linux Agent v198 also includes updates to improve agent security and performance.

- Windows Agent v155:
  - Corrected a rare memory leak in Windows event log monitoring when `LogMonitorUseServerTime` was set to `True`, rather than the default value of `False`. (Case: 00596915)
  - Windows 11 devices are now correctly reported by the agent. (Case: 00522253)
  - Adjusted summary file name timestamps to create the file name and timestamp before collection begins.
  - Improved agent debug log timestamps.
  - Updated the default value of the `DataKeep` parameter to 10 and limited the number of system logs included in diagnostic files .
  - Added a sync every three hours for Dynamic Applications aligned to the agent devices.
  - Windows Agent v155 also includes updates to improve agent security and performance.
- **For more information:** See the *Introduction to the Skylar One Agent* section.

Additional Agent Updates

- Added a new ***Gen3 Agent Unavailable Event Suppression Threshold*** field to the **Behavior Settings** page (System > Settings > Behavior). This setting limits the number of per-agent "unavailable" events during large outages to reduce event noise and improve triage during widespread agent outages. When the number of unavailable Gen 3 agents exceeds the configured threshold, the system generates a single, aggregated "Extended Outage" event instead of one event per agent, and then clears the event once availability recovers below the threshold.
- Updated the "ScienceLogic Agent" PowerPack to v103 with new internal events and alert definitions to support extended outage detection and recovery.
- Improved agent performance and resilience during data collection and export. A new health check now pauses the agent pipeline data export if the Extended Architecture cannot communicate with Skylar One. Gen 1 agents also support configurable collection timeouts to prevent failures with large data sets.
- Reduced bandwidth usage and resource consumption by enhancing agent upload behavior. Gen 1 agents now support summary uploads once per minute instead of every 20 seconds, with core agent details now published to Data Collectors for improved visibility. Additionally, you can now configure agents to upload data every five minutes instead of every minute. To do so, go to the **Agents** page (Devices > Agents), select an installed agent, then click the **[Settings]** tab and update the ***Upload Interval*** field to a value up to 300 seconds.
- Improved reliability and stability of core agent services, including performance and availability management components and the Agent Vitals service.
- Hardened agent startup, authentication, and error-handling behavior. Agents now shut down when provided an invalid token, recover more gracefully from transient back end (500-level) errors, and benefit from more robust Streamer Prime initialization logic.
- Improved Streamer Prime reliability and upgrade safety by isolating socket file paths and standardizing log handling through rsyslog, reducing permission-related failures during startup and upgrades.

- Fully removed support for Scylla from the Extended Architecture and agent pipeline. Scylla had previously been made optional and later deprecated in earlier release lines.

Device Management

- **What's new:** ***Enhanced topology settings.*** Added enhanced topology configuration options at the device level in the default user interface (AP2) to improve relationship modeling and topology accuracy. Four new fields are now available on the **[Settings]** tab of the **Device Investigator** page, allowing you to configure enhanced processing for:
    - Layer-2 topology
    - Layer-3 topology
    - Cisco Discovery Protocol (CDP)
    - Link Layer Discovery Protocol (LLDP)

  These options, which were previously available only in the classic user interface, enable Skylar One to form more accurate device relationships, generate richer topology maps, and improve event correlation by providing additional context about how devices are connected.

    - **For more information:** See the section on *The Settings Tab*.

Additional Device Management Updates

- Added the ability to merge or unmerge physical and component device records that represent the same device in the default user interface (AP2). From the **Devices** page, click the drop-down arrow next to the **[Add Devices]** button and select *Merge Devices* or *Unmerge Devices*.
- The field length for interface aliases in the `master_dev.device_interfaces` database table has been increased from 128 to 240 characters. The interface name displays in the ***Alias*** field on the **Interfaces** page (Registry > Networks > Interfaces) and on the **[Interfaces]** tab in **Device Properties** in the classic user interface.
- Discovery Dynamic Applications can now assign an IP address to the component devices they create. For more information, see the section on *Dynamic Applications Enhancements*.

Event Management

- **What's new:** ***A redesigned Event Suppressions page.*** Modernized the **Event Suppressions** page (Events > Suppressions) in the default user interface (AP2) to replace the classic, iframed version. This update improves usability and adds new management capabilities:
    - The updated **Event Suppressions** page allows you to view and unsuppress events directly from the page and displays detailed information including device or device group name, IP address, event policy details, severity, and creation metadata.
    - You can also remove one or more event suppressions from the updated **Event Suppressions** page.
    - The **Event Suppression List** page (Registry > Events > Suppressions) in the Skylar One classic user interface has been marked *deprecated* and will be removed in a future release.
    - **For more information:** See the section on *Responding to Events*.

Additional Event Management Updates

- Improved event filtering and table usability on the **Events** page and events-related dashboards:
  - Updated the visuals for the event severity filter at the top of the **Events** page so you can more easily see which severity levels are selected when filtering the events table.
  - Added column filtering for base custom attributes to the table on the **Events** page and to events-related dashboard widgets. This update supports a new filter syntax and partial matches for both string and integer attribute types.
- Enhanced event processing scalability and external ticket integration:
  - Improved event processing scalability and performance by enabling you to run multiple event processors in Skylar One simultaneously. You can configure the number of event processors by adding `eventing_num_executors =` to the `[LOCAL]` section of the **silo.conf** file and specifying a value between 1 and 75% of the CPU count.
  - Increased the time-to-live (TTL) for cached device group suppression in the event engine from one minute to five minutes to prevent long-running queries from getting stuck.
  - Added four new fields related to external tickets in the `master_events.events_active` and `master_events.events_cleared` database tables: `force_ticket_uri_2`, `force_ticket_uri_3`, `ext_ticket_ref_2`, and `ext_ticket_ref_3`.

Maps

- **What's new:** ***Geographic maps for business services and a new Locations page.*** You can now create interactive geographic maps for business services, enabling you to visualize your distributed business infrastructure across multiple sites, regions, and countries. When creating a geographic map, use the ***Entity Type*** field to choose whether to plot business services or devices. In addition:
  - The **Geographic Maps** page now displays a list of available geographic maps, including their entity type and the users who originally created or most recently edited the maps and the dates on which they did so.
  - A new **Locations** page (Manage > Locations) was added to Skylar One. From this page, you can add locations for devices or services, associate them with organizations, and assign locations to devices or services.

    NOTE: Version 102 of the "ScienceLogic: Geographic Maps" PowerPack includes a report that you can run to plot business services. A new bulk upload script for business services is also available on the *ScienceLogic Support Center*. ScienceLogic recommends using the report and script when initially setting up geographic maps for business services, and then using the **Locations** page for maintaining those locations.

  - Added new Access Hooks to manage user access for Geographic Maps and Location functionality. For more information, see *Authentication and Access Control*.
  - **For more information:** See the sections on *Geographic Maps* and *Locations*.

PowerPack Management

- **What's new:** ***A new PowerPack security and signing framework.*** This release introduces a new security framework for PowerPack management that enables digital signature verification and enforces trusted installation workflows. These changes help ensure the authenticity and integrity of PowerPacks, particularly in high-security environments such as in FedRAMP deployments. The following changes were made to support this enhancement:
  - Added a new ***Require Signed PowerPacks*** setting on the **Security Settings** page (System > Settings > Security) to implement system-wide controls for requiring signed PowerPacks. This setting is always enabled on Federal Information Processing Standards (FIPS) stacks and can be adjusted on non-FIPS stacks.
  - Defined the following PowerPack installation rules based on stack type:
    - On non-FIPS stacks, signed PowerPacks can always be installed if they have a trusted public key.
    - On non-FIPS stacks, unsigned PowerPacks can be installed only if the stack explicitly allows unsigned PowerPacks.
    - On FIPS stacks, only signed PowerPacks with trusted public keys can be installed.
  - Added a new **Trusted Public Keys** page (System > Manage > Public Keys) to the classic user interface that displays imported public keys and supports importing and bulk deletion of keys.

    NOTE: Only user-provided keys can be deleted.

  - Enhanced the PowerPack Installation workflow. The **PowerPack Installer** modal (System > Manage > PowerPacks > Actions > Install PowerPack) now includes a ***Signed*** column indicating whether a PowerPack is signed (*True*) or unsigned (*False*).
  - Updated the `systemSettings` GraphQL (GQL) resource to include the `requireSignedPowerPack` (boolean) field, and added the `updateRequireSignedPowerPack` mutation to manage this setting on non-FIPS stacks.
  - Added GNU Privacy Guard (GPG) public key management capabilities in GQL:
    - Created the `GpgPublicKey` resource with fields including `id`, `name`, `createUser`, `key`, and `restricted`.
    - Added `gpgPublicKey` and `gpgPublicKeys` queries, along with the `importGpgPublicKey` and `deleteGpgPublicKeys` mutations.

- Hardened PowerPack compatibility and encryption enforcement:
  - PowerPack operations now block unsupported encryption formats, returning "Unsupported PowerPack format" for PowerPacks exported from pre-11.3 systems.
  - You can no longer import or generate licenses for PowerPacks compiled prior to Skylar One version 11.1.
  - Internal calls from the PowerPack Engine now respect the global TLS-required setting when determining SSL usage.
- Created a formal process for installing out-of-band PowerPacks in FedRAMP environments, enabling secure verification and deployment using standardized manual commands or Access Hooks.
- **For more information:** See the *Introduction to Managing PowerPacks* section.

Skylar AI

- **What's new:** ***Improvements to the Anomaly Detection page.*** The following updates were made to the **Anomaly Detection** page to improve alert configuration workflows, modal behavior, and overall usability:
  - Clicking the **[Enable Alert Policies]** button now opens a modal that lets you select devices, Dynamic Application metrics, and indexes on which to enable alerting. The modal automatically filters out devices not aligned to organizations that use Skylar AI.
  - Enabling or disabling alerting now updates the alerting state while keeping the metric visible.
  - Ensured that the **Anomaly Detection Thresholds** page refreshes correctly without appearing to save changes when you click **[Cancel]**.
  - **For more information:** See the section on *Skylar Analytics: Anomaly Detection*.

Additional Skylar AI Updates

- Improved data export, ingestion, and metadata handling between Skylar One and Skylar AI:
    - The following data can now be exported from Skylar One for use in Skylar AI and Skylar Advisor:
        - Business services data, including owner organization ID; visible organization IDs; status; service policy name; and health, availability, and risk (HAR) refresh rate.
        - Event data, including event policy names, information about an event's cleared status, the name and ID of the user who acknowledged the event, and event-related ticket information.
        - Interface custom attributes
    - Optimized configuration data queries in the metadata exporter to improve performance.
    - Added new configuration options to the Skylar management script startup command:
        - `--verify-cert false`, which allows on-premises environments to connect using self-signed certificates
        - `--ca-bundle /path/to/bundle.pem`, which assigns a custom CA bundle to the `REQUESTS_CA_BUNDLE` environment variable
    - Added warning messages to debug logs for metadata exports when large device configuration tables may cause memory pressure.
    - Message payloads sent to Skylar AI are now automatically limited to less than 4MB, improving transfer reliability and network compatibility.
- Improved Skylar AI connection life cycle handling and indicator accuracy:
    - When a Skylar AI service connection is deleted, the Skylar management script now runs automatically to stop data export and reset feature toggles.
    - Updated how the Skylar AI ***Insights*** count is calculated on the **Device Investigator**. The value now reflects the total of:
        - Open root cause analysis events (accepted, suggested, or custom)
        - Open anomaly index score events
        - Open predictive alert events
    - The ***Insights*** count is now hidden entirely if the system is not connected to Skylar AI.

System Upgrade

- **What's new:** ***Significantly improved the speed, reliability, and visibility of the system upgrade and patching process.*** With this release, the following major updates were made to improve the system upgrade experience:
    - Enhanced patch import behavior so the patch import status is now marked as "Complete" immediately after the patch bundle is registered, which automatically starts the patch hook task.
    - Improved patch deployment performance and resiliency, particularly for larger Skylar One stacks, by revising the method in which deployments are batched and streamlining internal actions. This ensures faster and cleaner deployments, more resilient local deployment handling, and real-time progress visibility.
    - Added a new default timeout setting in the configuration for distributed appliances to improve the overall software update experience.
    - Adjusted system update staging rules so appliances that fail a patch hook are excluded only when running MariaDB 10.4.x, reducing unnecessary upgrade blocks on newer MariaDB versions.
    - Added detailed deployment-stage logs to the **System Updates** page (System > Tools > Updates) to improve visibility into the deployment process and support troubleshooting.
    - Updated the **System Update Logs** modal to highlight staging, pre-upgrade, and deployment entries based on defined severity rather than pattern matching, reducing false positives and clarifying success and failure states.
    - Shortened the time it takes to deploy updates by ensuring post-update scripts run only for packages that actually changed.
    - Added a post-update cleanup script that removes old patch bundles to free space on the Database Server. During upgrades, the system now retains only the two most recent patches and automatically removes all older patch bundles.

      > **NOTE:** This cleanup process can increase upgrade time during the first run, especially on systems with a large number of older patch bundles. Cleanup takes approximately 3–4 minutes per patch. For example, removing 15 older patches can add roughly 45–60 minutes. To reduce upgrade time, you can pre-clean older patches before upgrading by running `siloupdate clear patches` on the active Database Server.

    - **For more information:** See the section on *Updating Skylar One*.

Additional System Upgrade Updates

- Updated labels on the **System Updates** page (System > Tools > Updates) to use "Skylar One" instead of "SL1" or "EM7".

Topology

- **What's new:** ***Added support for Address Resolution Protocol (ARP)-based topology relationships.*** When ARP collection is enabled and ARP records are returned, Skylar One can automatically create device relationships without requiring full Layer-3 information from both endpoints. ARP relationship data can also be exported from Skylar One to Skylar AI for analytics and downstream use.
    - **For more information:** See the section on *ARP Topology*.

Additional Topology Updates

- Added enhanced topology configuration options at the device level in the default user interface (AP2) to improve relationship modeling and topology accuracy. For more information, see the section on *Device Management Enhancements*.
- Improved topology accuracy and relationship modeling:
    - Improved CDP and LLDP topology by allowing relationships to form when the interface is discovered only on the parent device, rather than both devices needing to be discovered, as long as enhanced CDP or LLDP processing is enabled.
    - Enhanced CDP relationship detection to discover additional node connections.
    - Stale Layer-2 topology relationships are now removed from maps when connections no longer exist or a parent/child device becomes inactive.
- Added the following enhanced topology configuration options on the **Behavior Settings** page (System > Settings > Behavior):
    - ***CDP & LLDP Enhanced Custom Attribute***, which lets you select an extended, string-type custom attribute to control device-matching order when enhanced topology processing is enabled.
    - ***L3: Ignore Hops for Undiscovered Devices***, which allows Layer-3 relationships to be created even when intermediate hops do not match a device in the system.
    - ***L3: Protocol Selection***, which lets you choose the traceroute protocol(s) used for Layer-3 topology (UDP, ICMP, or both).
- Expanded topology support across APIs and collection processes:
    - Added support for LLDP v2, enabling relationship discovery for devices that respond only to LLDP v2.
    - Devices in classes with PDU packing enabled now use bulk SNMP requests for Layer-2, CDP, and LLDP topology collection, drastically improving efficiency, reducing collection run times, and allowing more devices per Data Collector.
    - Added Layer-2, Layer-3, CDP, and LLDP topology settings to the REST API under the `device`, `device_class`, and `device_template` endpoints.
- Updated the "Enterprise Database: Topology Crunch" process to improve scaling.

Additional New Features and Enhancements in Skylar One 12.5.20

API

- Added comprehensive API key support across GraphQL and REST APIs. API keys now support authentication for GQL, can be used with REST API requests, include expiration metadata, and expose new `apiKey` and `apiKeys` GraphQL queries for management. API keys cannot be used to log in to the user interface.

  IMPORTANT: To use the new API keys feature, you must install the AP2 Quesito release after installing 12.5.20. The AP2 Quesito release includes a new **API Keys** management page that is required to use this feature.

- Added controls to ensure that SSH credential types created through the ScienceLogic API use private keys in the correct PEM format. In prior versions, that enforcement was only performed when the private key for an SSH credential was saved via the Skylar One user interface. For more information about the proper PEM format, see *Defining an SSH/Key Credential*.
- Expanded API support for event policy management, including the ability to create alert policies with embedded event definitions. The API supports severity assignments, occurrence logic, and automated associations between alert and event policies.
- Enhanced the Dynamic Application APIs to improve creation, retrieval, and performance. These updates include stricter schema validation and duplicate name checks when creating Dynamic Applications, support for retrieving snippet-based applications by ID or PowerPack, and response shaping via field exclusion to optimize payload size.

Authentication and Access Control

- Made the following updates to support the new API key feature:
  - Added a built-in user policy intended for read-only API access accounts, and ensured the Administration Portal access key requirement is enforced for both the default and classic user interfaces.
  - Added new database structures to support authentication and API key management, including the `master_access.api_key_meta` table and an `auth_method` column in the `master_account.accounts_deleted` table.

- Strengthened password handling and reset behavior with the following updates:
    - Removed several password- and authentication-related fields from the **User Policy Properties Editor** page (Registry > Accounts > User Policies > create/edit). Upon upgrading to Skylar One 12.5.20, these fields are reset to their default values for LDAP/Active Directory user policies. The following fields were removed:
        - ***Password Strength***
        - ***Password Expiration***
        - ***Password Shadowing***
        - ***Require Password Reset***
        - ***Authentication Method***
    - Local authentication users with MD5-hashed passwords are required to reset their passwords when upgrading to Skylar One 12.5.20.
    - Users cannot reuse their current password after expiration.
    - Password history rules are now enforced during password resets and updates.
    - Only user accounts configured for *Local* authentication can receive "I forgot my password" reset links.
    - On STIG systems, user accounts are automatically disabled after 35 days of inactivity and must be re-enabled by an administrator before the user can log in again.
- Standardized authentication configuration in the user interface:
    - When creating or editing a user account or user policy, you can now explicitly select *Local*, *AD/LDAP*, *SSO*, or *Certificate* in the ***Authentication Method*** field. This value is set automatically when a user logs in.
    - Improved consistency when managing authentication resources. When creating or editing an authentication resource:
        - Administrator user accounts can assign any user policy to the authentication resource.
        - Non-administrator user accounts can only assign non-administrator user policies that belong to an organization the user is a member of and includes only access keys the user has assigned to their user account.

      These controls apply both to the ***Policy*** drop-down field, which is filtered to show only accessible policies, and when you save.

      NOTE: Built-in user policies are not assigned to any organization. Because non-administrator users can only access policies whose organization matches one of their own, built-in policies will not appear in the ***Policy*** drop-down for non-administrator users and therefore cannot be assigned by them. To enable a non-administrator user to assign a built-in user policy, edit the policy so its organization is one the user belongs to.

- Added the following new Access Hooks to the "Platform - Administration" Access Key:
    - DEV_INTERFACES_PROPERTIES_EDIT, which enables users to set the interface collection frequency to one minute at the device, device template, and interface levels.
    - PLAT_INTERFACES_PROPERTIES_EDIT, which enables users to set the ***Initially Discovered Interface Poll Rate*** frequency to one minute system-wide on the **Behavior Settings** page (System > Settings > Behavior).
- Added a new "Maps" Access Hook category to manage user access to Geographic Maps functionality. This category includes the following new Access Hooks:
    - MAPS_GEOMAP_ACCESS_EMBEDDED, which enables users to view geographic maps embedded in other pages.
    - MAPS_GEOMAP_EDIT, which enables users to edit an existing geographic map.
    - MAPS_GEOMAP_DELETE, which enables users to delete an existing geographic map.
    - MAPS_GEOMAP_VIEW, which enables users to view the list of defined geographic maps on the **Geographic Maps** page.
    - MAPS_GEOMAP_CREATE, which enables users to create new geographic maps.
    - MAPS_GEOMAP_REG_PAGE, which enables users to view the Maps > Geographic Maps menu option.
- Added a new "Geographic Locations" Access Hook category to manage user access to Geographic Locations functionality. This category includes the following new Access Hooks:
    - GEOLOC_ALIGN, which enables users to align devices, device groups, and services to geographic locations.
    - GEOLOC_EDIT, which enables users to edit an existing geographic location.
    - GEOLOC_DELETE, which enables users to delete an existing geographic location.
    - GEOLOC_VIEW, which enables users to view the list of defined geographic locations on the **Geographic Locations** page.
    - GEOLOC_CREATE, which enables users to create a new geographic location.
    - GEOLOC_REG_PAGE, which enables users to view the Manage > Geographic Locations menu option.
- Introduced a new Access Hook, BIZ_SVC_INSIGHTS_VIEW, to control access to the upcoming **Service Insights** page, which will be added in a future release. This Access Hook is included in the "Grant All" and "IT Services - Administration" Access Keys by default.

Backup Management

- Newly installed systems now have their SMB backup options set to use SMB version 3.0 and encryption. The backup was refactored to include better logging and error handling. Subcommand output is now logged to the main backup log file. In addition, Skylar One now filters the credentials that are allowed based on the backup type selected.

Business Services

- You can now create geographic maps for business services. For more information, see the section on *Maps*.
- Made the following refinements to the **Service Investigator** page:
    - In the **Timeline** panel, the swim lanes in the **Skylar AI** and **Changes** sections now display the total number of events. Additionally, these two sections automatically expand or collapse based on configuration, subscription, and event availability.
    - When creating or editing service policy rules on the **Service Investigator** page, the preview graph now uses raw, hourly, or daily summarized data based on the selected timespan, providing more accurate preview results.
    - In the **Timeline** panel, the **[RestorePoint]** tab under the **Changes** heading is now **[Skylar Compliance]**.
    - In the **Events** panel, the **[Skylar AI]** tab is now **Skylar Advisor**.
- Added new data retention rules to prune topology performance and on-demand performance data after seven days.

Credentials and Discovery

- External credential services using CyberArk will now detect if a password change is in progress and enable Skylar One functionality to interrupt authentication.
- Added a new `allow_empty_passwords` setting to the `master.system_settings_core` database table. This system-wide setting enables you to determine if credentials created in Skylar One can have empty passwords.

Data Collection and Retention

- The "Data Collection: SNMP Collector" process and concurrent SNMP feature no longer run using Docker containers on Data Collectors. They now run as native Linux processes. With this update, STIG deployments can now use the concurrent SNMP feature.
- Added three new threshold slider fields on the **Data Retention Settings** page (System > Settings > Data Retention) for business services health, availability, and risk data retention:
    - ***Raw Business Services HAR Data***
    - ***Hourly Business Services HAR Data***
    - ***Daily Business Services HAR Data***
- Made optimizations to data pruning functionality for improved data maintenance efficiency.

- Improved resiliency, scalability, and error handling across data collection and ingestion pipelines. These updates help prevent overload conditions, reduce collection gaps, and provide more reliable behavior during failures:
    - Added a circuit breaker pattern for the Streamer Push service's Dead Letter Queues (DLQs) to pause data ingestion when DLQ thresholds are exceeded, preventing database overload during error conditions.
    - Refactored the delivery pipeline for the "Enterprise Database: Collector Config Push" process to improve scalability and reliability.
    - Updated the logging format and improved timeout handling for the "Enterprise Database: Collector Config Push" process.
    - Enhanced the resiliency of the data pull process by allowing you to configure a list of storage objects in the **silo.conf** file that should be automatically rejected or ignored.
- Improved Concurrent PowerShell collection by refreshing Kerberos tickets based on their actual expiration time instead of a fixed 8-hour interval. This prevents collection gaps when tickets have shorter lifetimes.
- Added storage enhancements, including enabling the `fstrim` systemd timer, swap compression cache, and reflinks (Copy-on-Write or CoW) to improve storage efficiency and performance.

Documentation

- Revamped the documentation site at *https://docs.sciencelogic.com*. This revamp introduces five new product-specific documentation sites:
    - Skylar One: *https://docs.sciencelogic.com/skylar_one/*
    - Skylar One PowerPacks: *https://docs.sciencelogic.com/powerpacks/*
    - Skylar AI: *https://docs.sciencelogic.com/skylar_ai/*
    - Skylar Automation: *https://docs.sciencelogic.com/skylar_automation/*
    - Skylar Compliance: *https://docs.sciencelogic.com/skylar_compliance/*
- Added a new **[SkyeDocs AI Search]** button to the top of every documentation page. This AI search agent can help you find information in a friendly way, using conversational English instead of search terms.

Dynamic Applications

- Added an *IP Address* option to the ***Component Identifiers*** field on the **[Collections]** tab for Dynamic Applications. This allows discovery Dynamic Applications to assign an IP address to the component devices they create and enables you to search for component devices using their IP addresses.

- Updated the Dynamic Application API to use the platform's ***Verify SSL for Collections by Default*** setting as the default for SSL verification during collection, instead of using a hardcoded value.

> **NOTE:** This change does not affect credential types that do not support SSL verification, such as classic PowerShell or SSH credentials, and does not affect SNMP.

- Modernized and enhanced management information base (MIB) compilation, import, and management workflows to improve reliability, scalability, and usability across the platform:
  - Updated the **MIB Compiler** page (System > Tools > MIB Compiler), converted the MIB compiler functionality from PHP to Python, and implemented the *pysmi* library to address issues with lost MIBs, missing symbolic name resolution for trap-based event policies, and related reliability problems.
  - Enhanced MIB import functionality to allow you to select and import multiple MIB files simultaneously.
  - When importing MIBs, you must click the **[Update]** button on the **OID Browser** page (System > Tools > OID Browser) to push updated MIB data to all Data Collectors in your stack. You must also click **[Update]** when bringing a new Data Collector online after additional MIBs have already been added. This manual sync ensures that all Data Collectors have the latest compiled MIBs, which are required by the event engine process when generating SNMP trap event details.
  - Improved the MIB Compiler by converting the recompile request to a POST method. This ensures that browser refreshes do not unintentionally trigger multiple redundant MIB recompilations.

EKMS

- Improved reliability and error handling for Enterprise Key Management System (EKMS)-based credential management to reduce failures during normal operation and system upgrades:
  - Added better exception handling when password errors occur in EKMS, allowing features that depend on vault credentials to handle failures more intelligently and display more informative error messages.
  - Modified how EKMS-based credential updates are pushed to distributed systems to improve reliability and consistency.
- Strengthened credential rollback behavior to reduce upgrade staging failures and ensure database consistency:
  - Added a rollback script to revert credentials in the `master.system_settings_licenses` database table away from EKMS-based encryption, reducing upgrade staging failures.
  - Ensured that encrypted credential rollback also updates credentials stored in the `master.system_settings_licenses` database table to maintain consistency.

- Enhanced the `slsctl` tool to improve compatibility and behavior across system roles and versions:
  - The `slsctl health_check` command is now skipped on passive nodes to prevent unnecessary or incorrect health evaluations.
  - Added backward compatibility to the `slsctl` tool so it can use configuration push methods across other supported Skylar One versions.

Global Manager

- Stack-level authentication is now enabled by default for Global Manager systems, and no longer requires you to manually add feature toggles during setup.
- To streamline access management across connected stacks, user accounts are now automatically created on a Global Manager stack when a request includes the necessary user and policy information. If the account does not already exist, it is created and the request continues as that user.

GraphQL (GQL)

- Expanded GQL support for enhanced topology and device relationship management:
  - Added enhanced-topology fields to the `updateDevice` mutation and the `devices` and `deviceClasses` queries: `l3_topo`, `lldp_topo`, `cdp_topo`, and `l2_topo`.
  - Added GQL support for AP2 device merging and unmerging, including `mergeDevices` and `unmergeDevices` mutations, the `mergedDevice` field on the `Device` entity, the `hasMergedDevice` search parameter, and `isMergablePhysicalDevice` and `isMergableComponenetDevice` boolean searches on the device query.
- Added new business services insights and adoption metrics in GQL:
  - Added the `businessServiceInsights` GQL query, which returns metrics including total services; counts of root, shared, device, disabled, maintenance, and orphaned services; device services with no associated devices; top 10 services by device count; policies using Dynamic Applications; and Dynamic Application metrics collected per run.

- Improved Global Manager ownership, global record control, and cross-stack visibility in GQL:
  - Updated `AccountPolicy` and related mutations to enforce ownership via globally unique identifier (GUID) headers/system IDs, support all fields in `updateAccountPolicy`, enable global management/propagation for `AccountPolicy` and `Organization`, and add `applyAccountPolicyToAccounts`.
  - Added claim/release workflows for global control, including an `accountPolicy` mutation to claim/release ownership, plus `claimAccountPolicy`, `claimAccessKey`, and `claimOrganization`. The `accountPolicies` query now includes `stackDiff` and `syncStatus` on Global Manager, and `ownerSystemId` is exposed read-only on globally managed queries.
  - Added Global Manager credential alignment improvements, including `createGlobalManagerStackCredential`, `alignGlobalManagerStackCredential` (replacing `alignGlobalManagerCredential`), and `unalignGlobalManagerStackCredential`.
  - Added `disableSkylarConnector` to remotely schedule the Skylar management script to be disabled.
  - Added experimental `GlobalDiff` and `GlobalView` queries for comparing a single global record across Global Manager and managed stacks for `account`, `accountPolicy`, `accessKey`, and `organization`.

    NOTE: These queries are available only on Global Manager systems and are designed for focused comparisons on a single global ID at a time to avoid performance issues.

- Tightened GQL schema consistency and compatibility controls:
  - Replaced `AccessKey.hooks` and the `hooks` argument on `createAccessKey` and `updateAccessKey` with `accessKeyHooks`.
  - Added the ability to filter GQL queries and mutations by minimum platform version at the field and argument level.
- Expanded GQL support for PowerPack governance and GPG key management:
  - Added `requireSignedPowerPack` to the `systemSettings` resource and the `updateRequireSignedPowerPack` mutation to control whether only signed PowerPacks can be installed (always true on FIPS systems).
  - Added the `GpgPublicKey` resource, `gpgPublicKey`/`gpgPublicKeys` queries, and `importGpgPublicKey`/`deleteGpgPublicKeys` mutations.
  - Added `uploadedPowerPacks` to list uploaded PowerPacks and view details such as filename, name, version, revision, and size.

- Improved reporting, credential, and administration support queries:
  - Added `classicReport` and `classicReports` to retrieve reports by ID, name, and PowerPack alignment, or to find reports not linked to any PowerPack.
  - Updated `credentialFields` to support searches for PowerPack alignment.
  - Added a parameter to `deleteGeoLocations` to allow deletion even when locations have aligned entities (devices or services); aligned entities are automatically unaligned upon deletion.
  - Added `getEventSuppressions` to support AP2 column filtering and delete actions, returning fields including `eventPolicy`, `dateCreated`, and `editedBy`.
  - Added GQL support for column filtering and sorting on the **Event Categories** page (replacing **Event Category Manager** in a future release).

PhoneHome Collector

- RSA256 and RSA512 algorithms can now be used for key authentication when configuring PhoneHome communication.

Platform and Security

- Skylar One version 12.5.20 includes package updates to improve security and system performance.
- Modernized the runtime, operating system, and database stack to improve compatibility and maintainability:
  - New ISO installations now use Oracle Linux 8.10 to support newer hardware platforms.
  - Updated MariaDB to version 10.11.16, aligned its configuration with system-wide encryption policies, and enhanced the security of internal MariaDB certificates.
  - The SQL schema is now managed as a dedicated package so schema checks and updates occur only when a new schema version is explicitly shipped.
  - Removed all Docker-related RPMs and dependencies from the non-extended Skylar One operating system.
  - Updated the Nginx stream module to version 1.24.
- Made the following changes to Python support:
  - Removed support for Python 2.7 execution environments. Skylar One now blocks the creation of Python 2.7 snippets and defaults to supported Python 3 versions.
  - Added support for Python 3.12. Bundled Python packages now support Python 3.9 through 3.12, and Python 3.12 is supported as a community runtime. After upgrading, execution environments that were built for Python 3.11 will be redeployed in Python 3.12.

- Improved orchestration, clustering, and update behavior:
  - Updated the template updater script to support the new Pacemaker configuration format on Oracle Linux 8, ensuring High Availability and Disaster Recovery clusters can be correctly patched and migrated.
  - The "EM7 Core: Task Manager" process now stops the `em7_scheduler` service when a Database Server becomes the passive node in a High Availability configuration.
  - Upgraded embedded Kubernetes components to version 1.35.
  - Updated Kafka to version 4.11 across AWS and on-premises deployment types.
  - Virtual environments now run from a compressed RAM disk; this change takes effect after reboot.
- Improved system security, usability, and troubleshooting:
  - Addressed a permissions issue that prevented vlock-protected terminals from being unlocked, and prevented TMUX from reusing SSH sessions that are already attached but locked.
  - Hid non-impactful error messages that previously appeared during the login process.
  - The `system_status.sh` script now includes the AP2 (NextUI) version in its output to aid support and troubleshooting.
  - The installation wizard now ensures the validity of FedRAMP ISO installations.
- Enforced `sudo` password requirements on STIG cloud deployments in AWS and Azure:
  - STIG systems deployed in the cloud now require a password when using `sudo` to run commands.
  - For new installations, you must set a password for the `em7admin` account before completing setup.
  - For newly installed and patched systems, a **Message of the Day** warning appears if no password is set. Run: `sudo /opt/em7/share/scripts/fix_sudo_passwd`

PowerPacks

- Updated and rebranded multiple PowerPacks to align with the Skylar One and ScienceLogic product naming:
    - Rebranded the "SL1 Default Dashboards Base Pack" as the "Skylar One Default Dashboards Base Pack" and updated all dashboard references from "SL1" to "Skylar One".
    - Rebranded the "SL1 Unified Theme" PowerPack to "Skylar One Unified Theme Base Pack", including updates to theme name and descriptive metadata.
    - Renamed additional PowerPacks to reflect updated ScienceLogic branding:
        - "EM7 Dashboard Widgets" is now "ScienceLogic: Dashboard Widgets"
        - "EM7 Default API Events is now "ScienceLogic: Default API Events"
        - "EM7 Default Credentials" is now "ScienceLogic: Credentials"
        - "EM7 Nmap Device Classes" is now "ScienceLogic: Nmap Device Classes"
        - "EM7 Standard Device Categories" is now "ScienceLogic: Standard Device Categories"
        - "EM7 Virtual Device Classes" is now "ScienceLogic: Virtual Device Classes"
        - "Interface Billing" is now "ScienceLogic: Interface Billing"
        - "ScienceLogic EM7 Base Pack" is now "ScienceLogic: Base Pack"
        - "SL1 Core Reports" is now "ScienceLogic: Core Reports"
        - "SL1 Credential Tests" is now "ScienceLogic: Credential Tests"
        - "SL1 Default Internal Events" is now "ScienceLogic: Default Internal Events"
        - "SL1 Performance Reports" is now "ScienceLogic: Performance Reports"
    - New installations no longer include the "EM7 Scheduled Dashboard Resources" PowerPack, because scheduled dashboard captures are no longer supported in the classic Skylar One user interface.

- Updated and enhanced several support and monitoring PowerPacks:
  - Updated the "ScienceLogic: Default Events" PowerPack to version 100.
  - Updated the "ScienceLogic: DRBD Monitoring" PowerPack to version 105.
  - Updated the "ScienceLogic: Support" PowerPack to version 110.

    IMPORTANT: ***Before updating the latter two PowerPacks***, you must remove all DRBD Dynamic Applications from the "ScienceLogic: Support" PowerPack and add them to the "ScienceLogic: DRBD Monitoring" PowerPack.

    Made the following updates in version 110 of the "ScienceLogic: Support" PowerPack:

    - Added new Dynamic Applications to monitor storage features, including ZSWAP, ZRAM, and LVM Thin Pools, with alerting for resource contention and high usage.
    - Added a new Dynamic Application to collect Non-Uniform Memory Access (NUMA) information and statistics. This Dynamic Application automatically aligns during discovery when NUMA hardware features are detected.
    - Added a new Dynamic Application to collect pressure stall information, which automatically aligns during discovery when supported.

Snippet Framework

- Improved execution behavior and performance for snippet-based collection:
  - Snippet-based collection no longer requires manual configuration of worker processes. The Data Collector now dynamically allocates worker processes based on available capacity and the jobs being executed.
  - Snippet Framework Configuration and Snippet Framework Performance Dynamic Application types now write a termination entry to the system logs when they terminate collection without completing (SIGTERM), including a list of pending device work items.
  - You can now download Dynamic Application execution logs, including failed runs, to support troubleshooting.

- Enhanced the snippet authoring, debugging, and inspection experience for snippet-based Dynamic Applications:
    - Introduced broad usability and troubleshooting improvements to the Snippet Framework user interface, including a modern code editor with line numbers, syntax highlighting, linting, and backend-sourced default snippet insertion with fallback behavior if the endpoint is unavailable.
    - The **Collection Object Editor** now emphasizes snippet argument authoring with **[Save]** and **[Save As]**, validation, device selection for test runs, optional properties and registry panels, and a run view that displays step-by-step execution output.
    - Added a collection object deduplication graph view that visualizes execution steps and how they are deduplicated. The graph can be downloaded as an SVG file.
    - Dynamic Application debugging views now include colored log highlighting to make warnings, errors, and other important events easier to identify.
    - The **[Properties]** view in the snippet-based editor was updated to align with modern styling, and the ***Dark Mode*** toggle was moved to the shared button bar for consistency across views.
    - The device inspection view now includes a table-based layout for device information, improving visibility into substitution values and expediting troubleshooting.
    - Added a live JMESPath expression evaluator to the Snippet Framework user interface, enabling inline testing of selectors against existing JSON payloads, with indexing controls and direct links to JMESPath documentation.
- Expanded tooling support, deployment flexibility, and execution environments:
    - Version 104 of the "Low-Code Tools" PowerPack is now included in the Skylar One 12.5.20 ISO, providing updated low-code Dynamic Application Builder capabilities out of the box.
    - The Snippet Framework user interface service is now enabled by default, reducing setup steps before first use.
    - Snippet-based Dynamic Application tooling can now run on Administration Portal appliances in distributed Skylar One deployments.
    - The third-party library download workflow can now automatically align downloaded wheels to the current execution environment, replacing existing libraries with a notification when applicable.
    - Added a new ***Python Wheel Downloader Pip Options (JSON)*** field to the **Behavior Settings** page (System > Settings > Behavior). This field enables you to specify custom pip repository URLs and trusted hosts for the Snippet Framework's Wheel Downloader. If you do not make updates to the field, the system preserves the default behavior.

- Strengthened governance, access control, and observability for low-code workflows:
    - Added organization-level audit logging for key Snippet Framework and Low-Code Explorer actions, including library management, wheel downloads, environment alignment, snippet linting, and collection runs.
    - Improved identity validation and access control for Low-Code Explorer endpoints. All endpoints now enforce user authentication (401) and permission checks (403) for snippet, library, environment, and execution actions.
    - Defined memory limits for the Low-Code Explorer service using configurable `MemoryHigh` and `MemoryMax` values to improve memory usage while preserving standard development and troubleshooting workflows.
    - You can now select multiple aligned devices when running multi-device actions, including generating execution graphs, instead of being limited to single-device views.
    - Updated access to the latest Dynamic Application API, ensuring availability of the most recent enhancements and stability fixes.

Subscription Billing

- Updated the `sl_subscription_usage_crunch` version in the **requirements.txt** file to add `ap2_version` to the `sys_info` object in the `sys_config` billing payload.

User Interface

- Introduced deprecation banners across selected pages in the classic user interface to signal features that will be removed in a future release. A banner now appears on the following classic user interface pages:
    - **Event Console** (the Events tab)
    - **Event Policy Manager** (Registry > Events > Event Manager)
    - **Event Suppressions** (Registry > Events > Suppressions)
    - **IT Service Manager** (Registry > IT Services)
    - **IT Service Dashboards** (Registry > IT Services > IT Service Dashboards)
    - **SLA Definitions** (Registry > IT Services > SLA Definitions)
    - All pages opened by clicking the **[Guide]** button

> **NOTE:** Content in classic user interface guides was last updated in SL1 8.12.2. For current documentation, use the **[Help]** button in the default user interface (AP2) or visit *https://docs.sciencelogic.com*.

    - Additionally, the following options on the **Account Preferences** page (Preferences > Account > Preferences) are now marked as deprecated:
        - Group by Organization
        - Show Masked Events
        - Collapse All Organizations
        - Event Console Columns pane
- Updated product branding throughout the user interface to reflect the transition from "EM7" and "SL1" to "Skylar One", including:
    - Home dashboard references
    - All former "EM7" device classes
    - ISO Installation window content
    - Node Configuration Utility content
    - Additional user interface text and labels
- Added a confirmation dialog when clearing or deleting a Dynamic Application from the **Dynamic Applications Manager** page (System > Manage > Dynamic Applications) in the classic user interface.
- Removed the *Schedule dashboard* option from the **Dashboards** page in the classic user interface, as scheduled dashboard captures are no longer supported.

Issues Addressed in Skylar One 12.5.20

This section describes the issues that were addressed in Skylar One 12.5.20.

Access Control

- Resolved an issue that prevented the ***Edited By*** column value on the **User Accounts** page (Registry > Accounts > User Accounts) from updating properly after an account-level change was made. (Cases: 00278274, 00515490) (Jira ID: EM-52869)
- Added a new DEV_REPORT_CREATOR Access Hook to ensure the **[Report Creator]** button (printer icon) and its report actions can be disabled on a per-user basis when required. (Case: 00453866) (Jira ID: EM-66141)
- Ensured that users can view records for their assigned organizations on the **Device Processes** (Devices > Processes) and **Windows Services** (Devices > Services) pages in the classic user interface. (Case: 00544862) (Jira ID: EM-76329)
- Ensured the **Access Keys** page updates automatically whenever a new access key is created. (Jira ID: EM-74431)

Agent

- Improved the resilience of Windows CPU and disk data collection so the aggregator service is less likely to run out of memory. (Case: 00616233) (Jira ID: EM-78721)
- Improved Gen 1 agent logic to reliably collect metric data after device merges. (Case: 00588555) (Jira ID: EM-78142)
- When deploying the Skylar One Extended Architecture, you can now use an asterisk (*) in the hostname for the ingress endpoints for the Responder and Streamer services. (Case: 00555223) (Jira ID: EM-76758)
- Ensured that the correct agent packages are included in the release. (Case: 00544019) (Jira ID: EM-76401)
- Addressed an issue that caused "HTTP Status Response 400" events when aligning Dynamic Applications to a large number of Gen 3 agent-monitored devices. (Cases: 00534237, 00546962, 00553238) (Jira ID: EM-76247)
- Resolved an issue that caused gaps in collection for agent-based devices due to process alert threshold validation in extended environments. (Case: 00514295) (Jira ID: EM-73890)
- Suppressed agent availability updates during extended cluster health issues to reduce alert noise and raise a system-level alert instead. (Cases: 00449747, 00450274) (Jira ID: EM-65707)
- The agent summary period is now set to a default of 60 seconds, with the option to change it to up to 300 seconds. (Cases: 00589019, 00589175) (Jira ID: EM-78064)
- Resolved an issue that prevented users from upgrading or deploying agents to stacks running Extended Architecture versions prior to 12.3.8. (Case: 00548279) (Jira ID: EM-76983)

- Addressed an issue where the **[Configs]** tab was not populated correctly for Windows agent-monitored Dynamic Applications; re-alignment is required after upgrade. (Case: 00484654) (Jira ID: EM-71649)
- The Streamer Collect service for Gen 1 agents now produces fewer log messages when no Gen 1 devices are aligned, reducing unnecessary log noise. (Case: 00507528) (Jira ID: EM-62777)
- Resolved an issue that prevented the **Agents** page from being accessed or displayed for some environments. (Cases: 00573304, 00586748, 00552775, 00564899) (Jira ID: SLUI-23209)
- Ensured the default agent Dynamic Applications can align to agent-monitored devices. (Jira ID: EM-78517)
- Addressed issues with Gen 1 agents reporting port status incorrectly due to IPv6-only reporting on Ubuntu devices. (Jira ID: EM-74303)

API

- Prevented deletion of organizations that are set as the primary organization in a user policy, avoiding orphaned references and memory issues. (Case: 00364076) (Jira ID: EM-59704)
- Corrected the device API so merged devices return the proper parent device for accurate reporting. (Case: 00563232) (Jira ID: EM-77101)

Authentication

- Pruned stale records from the `master_access.access_failed` table to prevent unbounded growth and reduce database load. (Case: 00425349) (Jira ID: EM-63882)
- Resolved an issue that prevented the ***Single Instance Login*** setting on the **Behavior Settings** page (System > Settings > Behavior) from working as intended for local, non-administrator accounts that utilize CAC or non-ASCII ADFS authentication. (Jira ID: SLS-1674)
- Addressed an issue that prevented users who were using a yum proxy from logging in to the Web Configuration Utility. (Case: 00515777) (Jira ID: EM-74126)
- Resolved an issue where attempting to update the Web Config Utility login passwords for Skylar One appliances resulted in error messages. (Cases: 00504492, 00505294) (Jira IDs: EM-72698, EM-72959)

Business Services

- Fixed an issue where editing different service policy rules that use the same Dynamic Application but different metrics would display the wrong metric label on the Y-axis of the graph. Now, the correct metric label is shown for each rule when editing. (Case: 00562795) (Jira ID: SLUI-23118)
- Addressed an issue that prevented HTTP headers from being added to custom universal credentials, allowing headers to be configured and saved successfully. (Case: 00603323) (Jira ID: SLUI-23493)
- Resolved an issue that caused the **Business Services** page to display an error page after upgrading to Skylar One version 12.3.9. (Case: 00593376) (Jira ID: SLUI-23194)
- Addressed an issue where services deleted prior to the AP2 Lokma release were still appearing in the `data_har` database table despite their removal. (Jira ID: SLUI-21757)

- Fixed an issue where severity chips for healthy events were not displaying when *Cleared Events* was selected from the **[Log Insights]** tab on the **Service Investigator** page. Now, cleared Skylar Automated RCA events correctly show severity chips for healthy events. (Jira ID: SLUI-22256)
- Addressed an issue where, for services with the ***RCA Options*** field enabled, removing a child service prevented Skylar One from computing health, availability, and risk values until the Service Topology Engine returned an updated topology, which occurred every 5 minutes by default. (Jira ID: SLUI-18853)

Collector Groups

- Ensured that you cannot delete a Data Collector if it is the only Data Collector assigned to a collector group (CUG) and it has devices aligned to it. (Jira ID: EM-79578)

Credential Management and Discovery

- Resolved an issue with SOAP/XML credentials that use cURL options that caused the system to sometimes not replace %D or %N variables in the URL with the appropriate IP or hostname value. (Case: 00574619) (Jira ID: EM-72620)
- You can now save ed25519 SSH private keys, including OpenSSH-formatted keys and those with lines longer than 64 characters, in SSH/Key credentials. (Case: 00605214) (Jira ID: EM-78511)
- The Skylar One Credential Gateway Service is now properly managed as a systemd service, allowing control via `systemctl` commands and ensuring automatic restarts after reboot. (Case: 00541087) (Jira ID: EM-76264)
- The **snmptrapd.conf** update process has been improved on Data Collectors and Message Collectors to avoid file locking issues, ensuring SNMPv3 trap authentication updates are applied reliably. (Case: 00414008) (Jira ID: EM-63206)
- SNMP v3 credentials with a ***Security Level*** of *Authentication Only* now successfully pass credential tests regardless of their ***Privacy Protocol*** setting, improving compatibility with devices configured for authentication only. (Cases: 00310315, 00381254) (Jira ID: EM-56900)
- Changing a device's default SNMP Read credential now updates all associated Dynamic Applications on all Skylar One Collectors promptly, rather than taking an hour or more and causing failed collections. (Case: 00565997) (Jira ID: EM-77443)
- Resolved an issue in the **Create Credential** modal that appeared when creating an AWS credential, where the modal incorrectly displayed the ***Enable FIPS Endpoint*** and ***Enable SSL Verification for the SL1 API*** toggles. (Case: 00523734) (Jira ID: SOL-29705)
- Resolved an issue where the **Credentials** page occasionally did not load or display all credentials as expected. (Case: 00551140) (Jira ID: SLUI-22673)
- When defining or editing a credential in the **Credential Tester** panel of the **Edit Credential** modal on the **Credentials** page, the ***Select Credential Test*** field displays only the credential tests relevant to the selected credential type. (Case: 00526225) (Jira ID: SLUI-22371, SLUI-22754)
- Resolved an issue that sometimes caused a "red bell" error on the **Credentials** page, indicating that the page could not load. (Cases: 00568449, 00568858) (Jira ID: EM-77246)
- Ensured that when adding or updating SSH credentials, changes are now saved correctly when a valid Privacy Enhanced Mail (PEM) key is entered. (Jira ID: SLUI-23324)

- Addressed an issue where some fields in the **Create Credential** modal for certain credential types did not display guidance text. (Jira ID: SLUI-23265)
- Resolved an issue with retrieving the universal credential that prevented the ability to run backups from the Disaster Recovery node. (Jira ID: EM-78993)
- The `slsctl health_check` is now skipped on passive nodes. (Jira IDs: SLS-1404)
- Ensured that you can discover Oracle devices with Native Network Encryption (NNE) enabled or using secure communication (TCPS). (Jira ID: EM-74110)
- Resolved an issue where discovery deduplication failed to detect existing IP addresses for merged devices, resulting in duplicate devices being created during scheduled discoveries. (Cases: 00632067, 00652539) (Jira ID: EM-79304)

Dashboards

- Addressed an issue where dashboard widgets returned duplicate timestamps and values for certain performance metrics, causing incorrect data visualizations. (Case: 00488371) (Jira ID: EM-71750)
- Ensured that heat maps display as expected when creating or editing **Interfaces** widgets. (Case: 00552693) (Jira ID: SLUI-22686)
- Updated the **Dashboards** page to ensure it remains stable and displays data correctly, even in dashboards with a high number of widgets using the chart visualization options. (Case: 00519023) (Jira ID: SLUI-22480)
- Resolved an issue where **File System** widgets using the *Line Chart* visualization would sometimes display default data labels in the chart legend rather than the correct file system labels. (Case: 00498710) (Jira ID: SLUI-21890)
- Ensured that multi-tag filters for **Interfaces** widgets work reliably. (Case: 00598090) (Jira ID: SLUI-23313)
- Resolved an issue where editing the ***Collection State*** column in **Interfaces** widgets could cause widget data to disappear. (Jira ID: SLUI-22777)
- Dashboards now retain their shared organization settings after being edited by any user, regardless of if they are an administrator or not. (Jira ID: SLUI-22691)
- Resolved an issue where toggling on ***This widget can drive other widgets*** in a **Device** widget with *Table* visualization caused selected checkboxes to blink in the **Preview** pane. (Jira ID: SLUI-22678)
- Resolved an issue where Global Manager dashboards could not be shared with organizations other than "System"; sharing now works for all organizations. (Jira ID: SLUI-22669)
- Addressed an issue where the ***Scale prefix*** field appeared twice when adding metrics and properties to a widget. (Jira ID: SLUI-22508)
- Ensured all relevant devices appear correctly in dashboard widgets and filters as intended. (Jira ID: SLUI-22960)
- Ensured the default "Server" dashboard no longer displays errors on the page or in the **Total Network Traffic** widget. (Jira ID: SLUI-21831)
- Ensured that you can resize individual custom attribute columns and reorder custom attributes using the drag-and-drop feature in **Events** widgets using the *Table* visualization. (Jira ID: SLUI-23233)

- When editing a **Devices** or **File Systems** widget using the *Number* visualization, toggling off ***Show the unit*** no longer exposes a nonfunctional ***Show the prefix*** toggle. (Jira ID: SLUI-23147)
- Resolved an issue in classic dashboards where filtering could leave **Custom Table** widgets stuck showing no data even when broader filters were selected. (Jira ID: EM-75612)
- Resolved an issue that caused device mismatches in dashboards due to dynamic device group rules not being fully pruned from the database. (Case: 00627806) (Jira ID: EM-79164)

Data Collection and Retention

- The **Performance Multi Object/Device Table** report now includes the ***Mounted On*** value for the "IBM: AIX Filesystem" Dynamic Application. (Case: 00507256) (Jira ID: EM-73338)
- Lowered the logging trace level in LLDP topology collection to resolve an "Illegal mix of collations" error that caused the Event Engine to crash. (Case: 00519994) (Jira ID: EM-74524)
- Resolved unhandled exceptions in SOAP/XML performance Dynamic Applications. (Case: 00501856) (Jira ID: EM-72767)
- Ensured filesystem inventory collection skips only devices with invalid SNMPv3 credentials instead of terminating the entire collection. (Case: 00513018) (Jira ID: EM-73887)
- Improved Class-based Quality of Service (CBQoS) inventory handling so only invalid objects are discarded instead of all collected data. (Case: 00523571) (Jira ID: EM-75171)
- Added a database index to CBQoS performance tables to resolve 504 gateway timeout errors when loading graphs in the user interface. (Case: 00564376) (Jira ID: EM-77626)
- Improved the CBQoS Stats storage object to improve performance at scale. This update reduces query counts by processing all CBQoS metrics for a policy within a single storage object. (Cases: 00491326, 00529258) (Jira IDs: EM-72124, EM-76603)
- Updated the CBQoS Stats storage object to resolve "Rows Behind" issues caused by high query counts when processing CBQoS metrics. (Case: 00521312) (Jira ID: EM-74595)
- Addressed a zero division error exception in CBQoS collection that occurred when an interface reported a port speed of zero. (Case: 00550203) (Jira ID: EM-76582)
- Updated CBQoS inventory collection to stop unnecessary collection attempts on disabled interfaces, which were resulting in issues on systems with a large number of interfaces. (Case: 00511842) (Jira ID: EM-73684)
- Reduced database churn by optimizing how SNMPv3 engine IDs are stored. (Case: 00605594) (Jira ID: EM-79104)
- Addressed an issue where a "ValueError" message occurred when a collection object was filtered down to zero results. The platform now handles empty substituted object ID (OID) lists without interrupting the collection process. (Case: 00568856) (Jira ID: EM-77600)
- Addressed an issue that caused data collection to fail for some SNMP Dynamic Application collection objects if they used certain advanced OID field features and PDU packing or concurrent SNMP collection was enabled. (Case: 00494992) (Jira ID: EM-72393)
- Improved logic to prevent false service monitoring alerts when internal collection Dynamic Applications (ICDA) data collection fails. (Case: 00517693) (Jira ID: EM-74564)

- Addressed an issue where system restart alerts were not raised for ICDA devices if the previous uptime value was unavailable. (Case: 00419150) (Jira ID: EM-73984)
- Ensured that you can successfully run a Dynamic Application by clicking its lightning bolt icon in Python 3.11 execution environments. (Case: 00596982) (Jira ID: EM-78277)
- Addressed an issue with interface bandwidth collection where enabling ***Interface Index Change Detection*** on a device class caused collections to stop. (Case: 00580438) (Jira ID: EM-77788)
- "Index Label" class collection objects now properly store the index value when concurrent SNMP collection is enabled. (Case: 00510290) (Jira ID: EM-73626)
- Resolved an issue that cause the "Data Collection: Interface Bandwidth" process to sometimes experience unhandled exceptions when upgrading from a Skylar One (SL1) version using Python 2.7 for interface collection to one using Python 3.6. (Case: 00500434) (Jira ID: EM-72827)
- The active() function in alert formulas for custom Dynamic Applications is now case-insensitive, ensuring healthy events trigger as expected regardless of index case. (Case: 00490521) (Jira ID: EM-71972)
- Added special character escaping to SNMP Dynamic Applications so formulas with special characters are evaluated correctly. (Case: 00521127) (Jira ID: EM-74936)
- Resolved an issue where `rows_behind.py` could hang or take an excessive amount of time on newly configured Data Collectors. (Case: 00518661) (Jira ID: EM-74481)
- Resolved an issue that caused a "Too Many Open Files" event and system log message on the Enterprise Database Collector Task Manager. (Case: 00613184) (Jira ID: EM-78660)
- The "EM7 Core: Hourly Maintenance" process now removes vanished devices that have reached their purge timeout, as intended. (Case: 00562754) (Jira ID: EM-77064)
- Resolved an issue where the daily maintenance process would halt the pruning of expired data after it encountered an exception. (Case: 00619712) (Jira ID: EM-79102)
- Ensured agent device availability is not impacted by orphan interface pruning operations during daily maintenance. (Case: 00595642) (Jira ID: EM-78153)
- Addressed an issue where the "Async Maintenance" process could fail to update custom attribute values containing special characters. (Case: 00467909) (Jira ID: EM-67137)
- Expired IT service logs are now pruned daily, preventing unbounded growth and avoiding `/tmp` filesystem issues and dashboard slowdowns. By default, these logs are pruned after 180 days, but you can configure this setting in the `it_service_log_retention` field in the `master.system_settings_core` database table. (Case: 00384468) (Jira ID: EM-61432)
- Resolved an issue that allowed the System Uptime OID form to be saved even when required fields were missing (null). (Jira ID: EM-77844)
- Updated Distributed Replicated Block Device (DRBD) Dynamic Applications to automatically align only if DRBD/clustering is set up on the Skylar One system. DRBD Dynamic Applications that are currently aligned to systems without DRBD set up will not automatically be removed, but they can be removed manually. (Jira ID: EM-63541)
- The system will now generate minor events for any errors that occur during specific types of data pruning during the Daily Maintenance process, rather only writing the errors in debug logs. (Case: 00393245) (Jira ID: EM-62000)
- Ensured the Daily Maintenance process properly prunes ad-hoc reports, deleted devices, and deleted Dynamic Applications to clean up unused data. (Jira ID: EM-74476)

Deployment and Configuration

- Corrected a service dependency issue where you could not log into the Web Configuration Utility (`:7700`) on fresh ISO deployments until the system was rebooted or the process was manually restarted. (Jira ID: EM-78175)
- Addressed an issue on STIG-compliant systems where the fapolicyd service interfered with the RPM database, causing deployment commands to become unresponsive. (Jira ID: EM-76568)

Device Groups

- Ensured device group rule selectors for null values on base custom attributes return devices with no value assigned. (Cases: 00425257, 00433276, 00592083) (Jira ID: EM-63886)
- Users who own shared device groups can no longer be deleted; such users are now suspended so their device groups remain available to other users. (Cases: 00350085, 00463304) (Jira ID: EM-61101)
- Resolved an issue where the API did not return devices dynamically aligned to device groups. (Case: 00517270) (Jira ID: EM-74445)
- Ensured that the "Update Device State" process completes successfully when storing the state of a device group, preventing unhandled exceptions. (Jira ID: EM-74930)
- Resolved an issue where only one page of device matches displayed on the **Device Groups** page when the group contained child subgroups. (Case: 00453572) (Jira ID: EM-66519)

Device Management

- Ensured that SNMP timeouts or other errors occurring while attempting to retrieve a device's sysObjectID do not result in the removal of the device class. (Case: 00450503) (Jira ID: EM-71048)
- Optimized queries and loading logic for the **Device Components** registry in both the default user interface (AP2) and classic user interface to avoid timeouts and long load times in large environments. (Case: 00542677) (Jira ID: EM-76355)
- Resolved an issue that resulted in orphaned component devices after their root devices were deleted. (Case: 00503030) (Jira ID: EM-72853)
- Added a new `device_group_name` filter to the device API endpoint and ensured that existing `device_group` and `organization` filters work as intended. (Case: 00505039) (Jira ID: EM-72774)
- Improved the hourly IP address cleanup task so it no longer causes load spikes on the Database Server. (Case: 00626066) (Jira ID: EM-79101)
- Resolved an issue where SNMP read and write credentials on component devices could be reset to a nonexistent credential ID when saving changes on the **[Properties]** tab. (Case: 00524768) (Jira ID: EM-75106)
- When creating a physical device manually with no SNMP credential and selecting *Ping | ICMP* as the device class, the device now defaults to ICMP for availability instead of SNMP. (Case: 00467477) (Jira ID: EM-67100)

- Updated the `smtp_auth_helper` script so SMTP test functions execute correctly and event emails can be sent. (Case: 00538592) (Jira ID: EM-76477)
- Resolved an issue where clicking **[Save As]** on a Dynamic Application failed if associated event policies had defined suppressions. (Case: 00514890) (Jira ID: EM-73983)
- Resolved an issue that resulted in a blank page if you clicked the printer icon on the **Device Processes** page (Devices > Processes) or **Windows Services** page (Devices > Services) . (Case: 00503774) (Jira ID: EM-73062)
- Fixed an issue where devices could not be aligned to any organization other than "System" when Global Manager mode was disabled. (Jira ID: SLUI-22664)
- Improved the **Anomaly Chart** modal on the **[Anomaly Detection]** tab of the **Devices** page by adding visual dividers and updating tooltips. The top chart shows only anomaly scores and the bottom chart displays metric values and expected ranges with units. (Jira ID: SLUI-22379)
- The **[Anomaly Detection]** tab no longer appears in the **Device Investigator** if the device is not collecting data for Skylar Analytics. (Jira ID: SLUI-23144)
- Ensured that device logs indicate how many times an event has repeated. (Jira ID: EM-74381)
- Addressed an issue where a vanished device was not purged after the retention period passed. (Jira ID: EM-74867)
- Ensured that, when you attempt to bulk merge devices in the user interface, devices appear based on applied filters as intended and the count of potential devices to be merged is accurate. (Jira ID: EM-72322)
- Added a check to see if the presentation ID exists when initializing Dynamic Applications that monitor device vitals to ensure you can add device vital panels to the **Device Investigator**. (Jira ID: EM-76034)
- Addressed an issue that caused the REST API to return only the count of device notes rather than their full content. (Case: 00664777) (Jira ID: EM-80115)
- Ensured that, when a script is used to bulk delete devices, the deletion completes successfully. (Jira ID: EM-72592)
- Ensured that when newly created base custom attributes are marked as unique, you cannot assign the same value for them to different devices. (Jira ID: SLUI-23430)

> **NOTE:** This change does not apply to existing custom attributes with non-unique values. To apply this change to existing custom attributes, go to the **Custom Attribute Manager** page (Manage > Custom Attributes) in the Skylar One classic user interface, change the index type, and then save the custom attribute.

Device Templates

- When selecting a collector group on the **[Config]** tab of a device template and applying the template to a device, the device is now properly aligned to the selected collector group as intended. (Cases: 00471246, 00482918, 00506448) (Jira ID: EM-70819)
- Dynamic Application sub-template fields for thresholds and roll-ups can now be updated correctly through the API. (Case: 00373284) (Jira ID: EM-60487)

Discovery

- System log messages now provide additional information about conflicting devices and IP addresses, making it easier to resolve discovery issues. (Case: 00559777) (Jira ID: EM-76933)
- Decreased the likelihood of deadlocks occurring in High Availability (HA) or Disaster Recovery (DR) deployments due to Dynamic Component Map (DCM) storage errors caused by medium-frequency data pull processes. (Case: 00503245) (Jira ID: EM-72962)
- Updated the "Discovery: Nightly Update" process to ensure it completes correctly and mitigates unhandled exceptions. (Case: 00480561) (Jira ID: EM-71440)
- The ability to run a global discovery is now disabled for Dynamic Applications whose ***Operational State*** is set to *Disabled*, preventing unintended alignments. (Case: 00552681) (Jira ID: EM-76858)
- The discovery session API now verifies the `edited_by` value in POST requests, ensuring the correct user is recorded for session edits. (Case: 00326460) (Jira ID: EM-57094)

EKMS

- Resolved an issue that caused the firstboot script to never complete in some Enterprise Key Management Service (EKMS) configurations if you re-ran firstboot on an already initialized system. (Jira ID: SLS-1958)
- The sl-vaultmngt service is now restarted automatically after an upgrade to prevent permission errors from occurring. (Jira ID: SLS-1772)
- Addressed an issue where the `slsctl credentials rotate` command did not update the encryption version in the `master.system_settings_licenses` database table. (Jira ID: SLS-1698)

Events

- Reduced false alerts and noise in the **Event Console** and downstream systems by resolving an intermittent issue where devices in dynamically suppressed groups still triggered filesystem events. (Case: 00522883) (Jira ID: EM-75290)
- Resolved an issue that could cause the Event Engine to fail if it encountered UTF characters in the ***yName*** (sub-entity name) field of active events. (Case: 00544491) (Jira ID: EM-76287)
- Ensured the correct ***yName*** is set when multi-match is enabled so event auto-clearing functions as expected. (Case: 00474473) (Jira ID: EM-71022)
- Prevented storage object errors by catching regular expression (RegEx) compilation errors in event policies; invalid RegEx patterns now result in a warning instead of a storage failure. (Cases: 00525229, 00610151) (Jira ID: EM-75585)
- Resolved an issue that caused the Event Engine to encounter unhandled exceptions during trap filtering if an SNMP Trap Filter contained a ***Host Filter*** longer than 64 characters. (Jira ID: EM-74036)
- Updated device IP address change detection rules so events include proper element details, reducing spurious "Device IP address change is detected" alerts. (Case: 00556599) (Jira ID: EM-76914)

- Suppressed "Process already running" errors from the "Enterprise Database: Config Push" process unless execution exceeds a configurable threshold, reducing unnecessary major events and improving signal quality. (Case: 00592138) (Jira ID: EM-78106)
- Corrected the event aggregation process so alert data is aggregated properly and dashboards reflect current event status. (Case: 00563211) (Jira ID: EM-77129)
- The ***Categorize events with an external system*** field in **[Advanced]** tab of the **Event Policy Editor** page (Events > Event Policies > create or edit) now accepts alphanumeric, punctuation, and special characters as intended. (Case: 00553372) (Jira ID: SLUI-22657)
- The **Event Policy Editor** now works correctly if you have event policy permissions, regardless of your Dynamic Application access. (Case: 00534604) (Jira ID: SLUI-22583)
- Fixed an issue where the **[Edit Note]** button on the **Events** page did not function when multiple events were selected for note editing in Global Manager systems. (Jira ID: SLUI-21131)
- Improved the visibility of the **[Cancel]** button on the **Event Policies** page when using dark mode, making it easier to see and interact with the button. (Jira ID: SLUI-22338)
- Addressed an issue that caused the Event Engine to crash due to a problematic character in the `in_internal.messages` database table. (Case: 00656764) (Jira ID: EM-79952)

Global Manager

- Non-administrator users can now query all Global Manager stacks regardless of their organization alignment. With this fix, users are able to make requests to any stack from the Global Manager system, even when they are not aligned to the same organization as the stack's device, though the aligned device will remain hidden. (Case: 00588571) (Jira ID: SLUI-23168)
- Resolved an issue that redirected users to child stacks via their IP address rather than the fully qualified domain name (FQDN) when clicking on a managed device hyperlink in Global Manager systems. (Case: 00571716) (Jira ID: SLUI-23073)
- Ensured duplicated dashboards display as intended in Global Manager mode. (Jira ID: SLUI-22784)

GraphQL (GQL)

- The `createEventPolicy` mutation now displays clear error messages when the source field is missing or invalid, allowing event policies to be created successfully through both the user interface and GQL. (Jira ID: SLUI-22550)
- The `CreateCredentialField` mutation now prevents duplicate credentials, and credentials can be created and edited correctly. (Jira ID: SLUI-22487)
- Addressed an issue that caused event policy update mutations to return old configuration values instead of the latest updated values in response. (Case: 0574610) (Jira ID: SLUI-22996)
- Ensured that the `subscriptionLicenseUsageByDevice` query returns reliable subscription license usage by device results for automation and reporting workflows. (Case: 00583363) (Jira ID: SLUI-23178)
- Resolved an issue where device path GQL queries returned a fetch error when using specific process-based filter syntax. (Case: 00336564) (Jira ID: SLUI-23244)

- Ensured that GQL queries for device assets return the same ***Asset Tag*** value shows in the asset's API data. (Case: 00484529) (Jira ID: SLUI-23199)

High Availability and Disaster Recovery

- Added support for three or more HA initializations of EKMS. (Jira ID: SLS-1341)

Inbound Messaging

- Improved handling of unused or unparsable email headers so inbound email messages using Oracle UTL_MAIL packages are processed correctly. (Case: 00604327) (Jira ID: EM-78456)
- Increased the character limit for the ***Authorized Email Domains*** field on the **Email Settings** page to 1,024 characters, allowing more domains to be added for multi-tenant environments. (Case: 00489246) (Jira ID: EM-72038)
- Resolved an issue where inbound email configuration settings were unintentionally removed during an upgrade, causing emails to bounce back to the sender. (Case: 00525407) (Jira ID: EM-77785)
- Ensured accurate system variable values are sent in SNMP trap notifications by resolving an encoding issue. (Case: 00554189) (Jira ID: EM-76768)
- Resolved an issue that prevented SNMP trap messages from displaying correctly when the payload sequence was not a valid hex string. (Case: 00534742) (Jira ID: EM-75898)
- Improved inbound email event policy matching so emails intended to clear events function as expected. (Case: 00509496) (Jira ID: EM-74441)
- Addressed an issue where the authorized email domain value was not saved correctly, preventing inbound email processing. (Case: 00509457) (Jira ID: EM-73347)
- The Event Engine now translates SMIv1-formatted SNMP traps to SMIv2, ensuring traps are recognized and processed correctly regardless of originating SNMP version. (Case: 00400440) (Jira ID: EM-62451)

Installation and Configuration

- Ensured that the **Add Bonding Interface** modal displays as intended when you click the **[+ Add Bonding Interface]** button in the **Node Configuration Application** (*<ip-address-of-appliance>*:7700/node-config). (Jira ID: SLUI-22285)

Internal Collections

- Resolved an issue where interface data was not collected for tunnel or other interfaces without a physical address, despite interface index change detection being enabled for the device class. (Case: 00595591) (Jira ID: EM-78360)
- File system inventory collection now skips individual devices with invalid SNMPv3 credentials instead of terminating the entire process, ensuring collection continues for remaining devices. (Case: 00513018) (Jira ID: EM-73887)
- Ensured failed or incomplete OS Process monitoring collection does not result in false-positive alerts. (Case: 00485271) (Jira ID: EM-64198)

ITSM

- Resolved unhandled exceptions in classic IT Service management caused by missing `app_id` values when using hardware metrics with virtual devices. (Cases: 00515724, 00515952) (Jira IDs: EM-73994, EM-74055)

Maps

- Addressed an issue where the ***Organization*** field in the Geographic Maps creation page displayed an error message requiring an organization to be selected, even when one had already been chosen. (Jira ID: SLUI-22349)

MariaDB

- Filtered out unneeded warnings from MariaDB, reducing noise in log files and preventing spurious events during system boot. (Case: 00524110) (Jira ID: EM-75219)
- Resolved an issue that caused very high CPU usage, SIGTERMs, and collection failures on Data Collectors due to the same query running repeatedly. (Cases: 00502085, 00502942) (Jira ID: EM-73066)
- Updated the `master_log.notifier_log` database table to use `BigInt` as the data type for the ***ID*** column, preventing failures when viewing automation logs due to integer overflow. (Case: 00375841) (Jira ID: EM-61138)

Monitoring Policies

- In Web Content monitoring policies, you can now add "#no-transfer-encoding" to the end of the policy's URL to ensure that URLs reachable with HTTP 200 do not generate erroneous 501 errors due to an invalid Content-transfer-encoding header, thus preventing false content verifier alerts. (Case: 00543092) (Jira ID: EM-76675)
- The ***HTTP Auth Username:Password*** and ***Proxy Username:Password*** authentication fields are no longer returned in the API results for Web Content monitoring policies. However, you can still add or update these fields using the API as you did previously. In addition, a new ***Clear saved credential*** checkbox was added to both of those fields on the **Web Content Policy** editor modal. Selecting these checkboxes clears any credentials saved for those fields. If a credential exists for either field, it now displays as a series of asterisks. (Case: 00646372) (Jira ID: EM-79598)
- Optimized SQL queries used to populate the **Device Services** page in the classic user interface to prevent 504 gateway timeouts. (Case: 00503216) (Jira ID: EM-72987)
- Updated Web Content monitor debug output to avoid exposing HTTP authentication credentials while preserving required diagnostic detail. (Case: 00577929) (Jira ID: EM-77704)
- Resolved an issue that prevented Web Content monitoring policies from decoding content when encoding could not be detected from response headers or body content, which previously caused expression matching errors. (Case: 00534734) (Jira ID: EM-75892)

Platform and Security

- Resolved an issue where DHCPv6 addressing failed on appliances due to missing local firewall rules. (Case: 00546503) (Jira ID: EM-76464)
- Resolved an issue where backup status was not properly detected as running, which incorrectly triggered DRBD disconnection events. (Case: 00603287) (Jira ID: EM-78338)
- Resolved an issue that caused services to experience an approximately 15-minute delay before restarting on the High Availability failover node after a hard stop of the primary node. (Cases: 00459541, 00522323) (Jira IDs: EM-74711, EM-66492)
- Updated MySQL log rotation to ensure logs are flushed and compressed correctly, preventing the `/var/log` directory from filling up and impacting system stability. (Case: 00474964) (Jira ID: EM-71034)
- Updated the `system_status.sh` utility to stop starting or restarting the `siloupdate-spool` service, which is no longer required for system update functionality. (Case: 00451070) (Jira ID: EM-65832)
- Collector availability checks now honor collector group alert suppression settings on secondary databases. (Case: 00405513) (Jira ID: EM-62553)
- Improved metrics server diagnostics so the query script detects when it is run on a passive data engine and exits with guidance to run on the active engine instead, preventing misleading "no data found" errors. (Case: 00614015) (Jira ID: EM-78794)
- The password for the em7admin account must now be set to use `sudo`. (Jira ID: EM-73675)
- Updated Skylar One to use Network Manager to control and configure `dnsmasq` for DNS management. (Jira ID: EM-77623)
- Reduced interface data export message size to improve delivery reliability while preserving successful export and storage of Gen 3 agent interface metrics. (Case: 00613904) (Jira ID: EM-78679)
- Updated the self-monitoring scripts to generate fewer errors and other messages when those features are not available. (Jira ID: EM-74219)
- Updated the "Collector Outage" event policy in the "ScienceLogic: Default Internal Events" PowerPack to include more descriptive information about the policy and how to use it. (Jira ID: EM-58676)

PowerPacks

- In the "ScienceLogic Support Pack" PowerPack:
    - Updated the "Support: Configuration File Validation" Dynamic Application to ensure configuration checks are handled correctly when fields appear in a different order, preventing false alerts. (Cases: 00595076, 00509442) (Jira ID: EM-73594)
    - Updated the "Support: Appliance Validation" Dynamic Application to handle missing appliances and incomplete data more gracefully. (Jira ID: EM-54148)
- Resolved an issue on the **PowerPack Manager** page (System > Monitoring > PowerPacks) that prevented filtering Dynamic Applications by name while editing a PowerPack. (Jira ID: EM-79219)

- If a PowerPack installation fails, you are now prompted to contact ScienceLogic Support rather than the page becoming unresponsive. (Jira ID: EM-73677)

Relationships

- Resolved an issue that caused user-created ad hoc device topology relationships to be inadvertently pruned during execution of the "Enterprise Database: Topology Crunch" process. (Jira ID: EM-78702)
- Ensured that substitutions from `EM7_VALUES['%_parent_id']` in Run Book snippet actions correctly show the current parent device ID for component devices, even if the parent device changes after the relationship is created. (Case: 00323280) (Jira ID: EM-56300)

Reporting

- Ensured that file names are properly aligned with user validation rules so reports download successfully from the **Report Output Template** page. (Case: 00544908) (Jira ID: EM-76453)
- Reports that include images or charts generate properly in ODS format as intended. (Jira ID: EM-75401)
- Ensured that non-administrator users that have all of the Reporting access keys aligned to their user account can create a new schedule or see the archived reports. (Jira ID: EM-72259)
- Ensured that the ScienceLogic logo properly displays across all reports included with the "ScienceLogic: Core Reports" PowerPack. (Jira ID: EM-60416)
- Resolved an issue where a reserved keyword introduced in recent MySQL versions prevented the "Asset List" report from generating output. (Case: 00494179) (Jira ID: EM-72433)
- Resolved a unit conversion issue in the "Asset List" report that caused RAM values to display incorrectly. (Case: 00578340) (Jira ID: EM-77865)
- Eliminated duplicate rows in the "Device Availability (Totals Only)" report to ensure totals are generated once per selection. (Case: 00325215) (Jira ID: EM-57072)
- All information from the "Device Combo" report now appears at the top of the generated report output as intended. (Jira ID: EM-74144)
- Expanded file system threshold handling in "Device Thresholds" reports so device-specific Warning/Major and Critical overrides are reflected correctly. (Case: 00318407) (Jira ID: EM-56191)
- Ensured that the "Device Utilization" report accurately reflects CPU and memory usage values collected from devices. (Cases: 00445507, 00526684) (Jira ID: EM-65530)
- Updated "Event Detections" report versioning to prevent upgraded stacks from displaying the same report version with different output behavior, such as numeric versus text-based severity values. (Case: 00403068) (Jira ID: EM-63044)
- Improved interface average utilization calculations so percentages reported in "Interface Usage" reports are accurate. (Case: 00464211) (Jira ID: EM-67096)
- Corrected "Interface Usage" report averaging logic so utilization percentages no longer exceed 100 percent, resolving inconsistencies between reports and performance graphs. (Cases: 00379494, 00401898, 00406457) (Jira ID: EM-61297)

- Ensured that interface aliases are included when exporting "Interface Usage" reports. (Cases: 00520575, 00525158) (Jira ID: EM-79127)
- Resolved an issue where the "Interfaces In Use" report generated empty output when the ***Interfaces currently down only*** option was selected. (Case: 00356292) (Jira ID: EM-59142)

Run Book Automation

- Addressed performance issues related to run book automation processing in Extended Architecture deployments. (Case: 00617113) (Jira ID: EM-78837)
- The run book automation engine now stops creating automation tasks when the task queue is full, improving stability and preventing performance issues when a large number of tasks are added. (Case: 00218151) (Jira ID: EM-48764)

Schedule Management

- Disabling a maintenance schedule now automatically removes affected devices from maintenance mode, even in complex scheduling scenarios. (Case: 00543914) (Jira ID: EM-77156)
- Addressed an issue where non-recurring schedules created via the REST API did not display an end date in the **Schedule Manager** page (Preferences > Account > Schedule) when used with ServiceNow change requests. (Case: 00609709) (Jira ID: EM-78690)

Skylar AI

- Ensured Skylar AI features remain enabled after a High Availability failover event without manual intervention. (Jira ID: EM-78849)
- Optimized metadata exporter queries against large Dynamic Application tables, significantly reducing export times and database load. (Case: 00623628) (Jira ID: EM-79051)

Subscription Billing

- Addressed an issue where device license usage data was intermittently missing from reports and the user interface. The subscription crunch process was updated to prevent duplicate key errors from causing silent data loss during database updates. (Case: 00662281) (Jira ID: PTEL-2558)

System Administration

- Corrected an issue where the Admin Notifier system warning banner could continue to display after the underlying events were cleared, even when no active events appeared in the **Event Console**. (Case: 00426805) (Jira ID: SLUI-23393)

System Update

- Addressed an issue where the patch hook failed to update on remote appliances when concurrent patch import was enabled, which could later result in failed deployments. (Case: 00641315) (Jira ID: EM-79475)
- Resolved an issue where retrying previously failed RPM upgrades could produce false positives and cause deployment failures during system updates. (Case: 00544453) (Jira ID: EM-76308)
- The deployment process now correctly identifies disaster recovery nodes regardless of naming conventions, ensuring successful deployment. (Case: 00541487) (Jira ID: EM-76190)
- Ensured the MariaDB service is fully stopped before packages are upgraded, preventing crashes during system updates. (Case: 00516395) (Jira ID: EM-74025)
- Improved responsiveness of PhoneHome communication by enabling TCP keepalive during database health checks. (Cases: 00437135, 00513519) (Jira ID: EM-65069)
- Added logging for the `locate-update` command and the affected packages to help troubleshoot staging failures on remote appliances during system updates. (Case: 00419470) (Jira ID: EM-63452)
- Resolved an issue where appliances initially marked as unavailable would still be incorrectly triggered for deployment after a database upgrade. (Jira ID: EM-78477)
- Addressed an issue that caused deployment to sometimes stall during a batched system update due to the system reporting conflicting deployment status. (Jira ID: EM-74366)
- Ensured the package installation now correctly identifies and maintains the appropriate platform version, preventing repository mismatches and post-update failures. (Jira ID: EM-79203)
- Resolved an issue that caused "no space left on device" errors when importing a patch bundle with *`use_concurrent_import`* enabled. (Jira ID: EM-79357)
- Adjusted the system update staging rules so appliances that fail patch hook are excluded only when running MariaDB 10.4.x. This reduces the chances for upgrades being blocked and the need for workaround measures on appliances running newer versions of MariaDB. (Jira ID: EM-78116)

Ticketing

- Resolved multiple issues in the **Ticket Editor** page (Tickets > wrench icon). Fixes include restoring double-click functionality to open tickets, enabling scrolling for multi-line notes, and restoring the maximize button in the notes section. (Cases: 00601480, 00601484) (Jira IDs: EM-78351, EM-78352)

Topology

- Improved Layer-2, CDP, and LLDP topology collection efficiency, enabling more devices to be aligned to a single Data Collector. (Case: 00515516) (Jira ID: EM-73998)
- Ensured that CDP topology relationships are created only when the parent device interface is discovered and enhanced CDP processing is enabled for the parent device or device class. (Case: 00590470) (Jira ID: EM-78101)

- Resolved an issue that caused unhandled exceptions during Layer 3 topology data collection when the system encountered a device without a primary IP address. (Case: 00473950) (Jira ID: EM-71630)
- Ensured that LLDP topology collection can create multiple relationships between the same devices when appropriate discovery data is available. (Cases: 00444042, 00549735) (Jira ID: EM-6502)
- Prevented the creation of database records for interfaces that are discovered during topology processing but do not exist in the system. (Cases: 00377649, 00505904) (Jira ID: EM-61146)
- Resolved an issue that prevented LLDP topology relationships from being created between device interfaces with matching port numbers. (Case: 00633557) (Jira ID: EM-79286)
- Ensured that LLDP relationships are created only when interface information actually matches between connected devices. (Cases: 00607249, 00624013, 00651952) (Jira IDs: EM-78532, EM-79658)
- CDP topology collection now continues processing other devices when encountering malformed interface names instead of terminating abruptly. (Case: 00513024) (Jira ID: EM-73828)
- Layer-2 relationships involving multiple interfaces across the two devices now appear in the classic user interface on the **[Relationships]** tab of the **Device Properties** pane and on the **Device Relationships** page (Registry > Networks> Device Relationships). (Case: 00645969) (Jira ID: EM-79546)

User Interface

- The ***Device Class/Sub-class*** filter on the **Vanished Devices** page (Devices > Vanished Devices) now correctly supports the "!" (NOT) operator, ensuring accurate filtering results. (Case: 00497858) (Jira ID: EM-72538)
- The Dynamic Application alert editor now prevents users from saving alerts with invalid formulas, ensuring syntax errors are caught during configuration instead of at runtime. (Case: 00550690) (Jira ID: EM-76512)
- Ensured that, when adding a new access key on the **Key/Hook Alignment Editor** page (System > Manage > Access Keys > Key Manager), the form resets if you click the **[New]**button. (Jira ID: EM-76407)
- Ensured devices can be deleted successfully from the default user interface (AP2). (Cases: 00515983, 00546954) (Jira ID: SLUI-22280)
- Accounts with IP restrictions can now access the default user interface (AP2) as well as the classic user interface. (Cases: 00363146, 00371129, 00474483) (Jira ID: EM-59581)
- The **Agents** page (Devices > Agents) page now loads correctly when vanished agent-monitored devices exist. (Case: 00493976) (Jira ID: EM-72366)
- Resolved an issue where saving changes to a Data Collector on the **Appliances** page overwrote the database password, causing the Data Collector to become unavailable. (Case: 00529031) (Jira ID: EM-75689)
- Corrected mismatches between CBQoS table and graph values, ensuring performance graphs accurately reflect maximum values. (Case: 00547468) (Jira ID: EM-76813)

- Addressed an issue where, when running Skylar One on a computer with limited hardware resources, using the **Bandwidth Billing Editor** page (Registry > Service Provider Utilities > Bandwidth Billing > create/edit) would cause some computers to experience an infinite redirect loop and lock up. (Case: 00494962) (Jira ID: EM-72445)
- Ensured custom themes and logos display as intended on classic user interface pages. (Case: 00503523) (Jira ID: EM-72921)
- Added pagination to the **[SVC Policies]** tab when creating or editing device templates. (Case: 00334322) (Jira ID: EM-56815)
- Resolved a unit type mismatch in the file system performance statistics where mouse-over values were displaying incorrect labels and calculations. The graph mouse-over behavior now correctly aligns with the stored binary-based data. (Case: 00626073) (Jira ID: EM-79157)
- Optimized the SQL queries used to populate the **Device Processes** page (Devices > Processes) to prevent 504 gateway timeouts. (Cases: 00576200, 00511384) (Jira ID: EM-73682)
- Ensured the ***Runtime Offset*** column on the **Admin Processes** page (System > Settings > Admin Processes) displays values as intended. (Case: 00506663) (Jira ID: EM-73218)
- Links to the deprecated Classic Global Manager will no longer be visible in the classic user interface when the "ScienceLogic: Global Manager" PowerPack is installed. (Jira ID: EM-73995)

Windows Monitoring

- Added a customizable timeout setting to ensure that long-running concurrent PowerShell requests are closed out as intended. This change affects the default Python module used to monitor Windows Devices, "pyWinRm", as well as legacy PowerShell. (Case: 00586338) (Jira ID: EM-77988)

> **NOTE:** To implement this setting, use the database statement `INSERT master.system_custom_config (field, field_value) VALUES ('async_ps_cmd_timeout', <value>);` on the Database Server, where `<value>` is replaced by the timeout value, in seconds. The value can be between 30 and 300. The default setting is 210 seconds.
>
> This setting can also be customized per collector group by specifying the collector group ID in the statement `INSERT master.system_custom_config (field, field_value, cug_filter) VALUES ('async_ps_cmd_timeout_CUG<ID>', <value>, <ID>);` where `<ID>` is replaced by the collector group ID number and `<value>` is replaced by the timeout value, in seconds.

- The system no longer requires DNS resolution for local Windows accounts, ensuring PowerShell collection works as intended on servers using local accounts. (Case: 00557117) (Jira ID: EM-76871)
- Implemented a worker process shutdown notification mechanism for the PowerShell collector to prevent stalled collections and restore stable operation across multiple Data Collectors. (Case: 00586338) (Jira ID: EM-78145)
- Addressed delayed and failing concurrent PowerShell collection results after upgrade. (Case: 00586338) (Jira ID: EM-78223)
- Added a configurable timeout for long-running concurrent PowerShell requests. (Case: 00586338) (Jira ID: EM-77988)

- Resolved an issue where concurrent PowerShell collections failed when using CyberArk-managed credentials, ensuring collections run reliably with CyberArk integration. (Case: 00598480) (Jira ID: EM-78224)
- Addressed an issue that caused agentless Windows devices to sometimes stop reporting PowerShell performance data due to internal server errors in the Collector Pipeline. (Case: 00560303) (Jira ID: EM-77066)
- Enhanced error handling for WMI monitoring to ensure that connection failures and query processing errors are logged to the device log and visible in the user interface. Previously, these failures were only recorded in backend log files. This update lets you create event policies that can turn the device logs into events. (Case: 00517090) (Jira ID: EM-74570)
- Resolved issues handling Windows account names that begin with a dash, such as "-em7admin". (Jira ID: EM-78571)
- Prevented duplicate Windows service monitors when names differ only by case. (Jira ID: EM-72961)
- Resolved an issue that prevented you from monitoring Windows devices with IPv6 addresses using WMI. (Jira ID: EM-73384)

Installing and Upgrading Skylar One

IMPORTANT: You can consume Skylar One 12.5.20 only if you are upgrading from an earlier Skylar One version that *supports upgrades to this release*. There is no ISO version for version 12.5.20.

For a detailed overview of Skylar One, see the *Introduction to Skylar One* manual.

For detailed instructions on performing a new installation of Skylar One, see the *Installation and Initial Configuration* manual.

For detailed instructions on upgrading Skylar One, see the section on *Updating Skylar One* in the ***System Administration*** manual and the upgrade notes that are included in this document.

NOTE: ScienceLogic strongly recommends that you review the *Known Issues* for Skylar One at https://support.sciencelogic.com/s/known-issues#sort=relevancy before installing a new update.

For known issues specific to this release, see the *Known Issues* section of this document.

Skylar One Extended Architecture

For existing on-premises deployments of Skylar One Extended Architecture, see the section on *Upgrading Skylar One Extended Architecture* in the ***System Administration*** manual for upgrade instructions. For help with technical issues, contact ScienceLogic Customer Support.

NOTE: New installations of Skylar One Extended Architecture are available only on SaaS deployments.

Important Upgrade Notes for Skylar One 12.5.20

This section includes important notes for upgrading existing Skylar One systems to the 12.5.20 release.

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

CAUTION: ScienceLogic strongly recommends that you review these notes in their entirety before upgrading to version 12.5.20.

Supported Upgrade Paths

The ScienceLogic-validated upgrade paths to version 12.5.20 are outlined below.

- 12.5.7 > 12.5.20
- 12.5.5 > 12.5.20
- 12.5.4 > 12.5.20
- 12.3.14 > 12.5.20

NOTE: For more information about supported upgrade paths for Skylar One releases, see the *Skylar One Recommended Upgrade Paths* section in the ***System Administration*** manual.

Unsupported Upgrade Paths

The following upgrade path is not supported due to release timing:

- 12.5.8 > 12.5.20

Pre-Import Script for Upgrade Efficiency

Beginning with 12.5.1, the system update import experience was improved to make it faster and more efficient by allowing you to run concurrent imports. This feature is not enabled by default. To enable it, download the patch bundle you want to import, then run the following script:

```
silo_mysql -e "INSERT INTO master.system_settings_patcher (param, value,
description) VALUES ('use_concurrent_import', 1, 'Use concurrent patch
import');"
```

After running the script, you can start the import process.

STIG Support

12.2.x and 12.3.x STIG-compliant users should contact their ScienceLogic account managers for information about upgrading to this release.

> **NOTE:** When deploying a STIG-compliant configuration, port 7700, the Web Configuration Utility, and the **Database Tool** page are all disabled. In addition, concurrent network interface collection is not supported for these deployments.

Upgrading MariaDB and Rebooting Skylar One

All Skylar One versions include important package security updates. To apply these updates, you must upgrade MariaDB and then reboot all Skylar One appliances. The following table specifies the required MariaDB version for each Skylar One version:

Skylar One Release	Release Type	Required MariaDB Version
12.5.20	ISO and upgrade	10.11.16
12.5.8	Upgrade only	10.6.21
12.5.7	Upgrade only	10.6.21
12.5.6	Upgrade only	10.6.21
12.5.5	Upgrade only	10.6.21
12.5.4	ISO and upgrade	10.6.21
12.5.3	Upgrade only	10.6.21
12.5.2	Upgrade only	10.6.21
12.5.1	ISO and upgrade	10.6.21
12.3.14	Upgrade only	10.6.18
12.3.13	Upgrade only	10.6.18
12.3.12	Upgrade only	10.6.18
12.3.11	ISO and upgrade	10.6.18
12.3.10	Upgrade only	10.6.18
12.3.9	Upgrade only	10.6.18
12.3.8	Upgrade only	10.6.18
12.3.7	ISO and upgrade	10.6.18
12.3.6	Upgrade only	10.6.18
12.3.5	Upgrade only	10.6.18
12.3.4	Upgrade only	10.6.18
12.3.3	Upgrade only	10.6.18

Skylar One Release	Release Type	Required MariaDB Version
12.3.2	Upgrade only	10.6.18
12.3.1	ISO and upgrade	10.6.18
12.3.0	ISO and upgrade	10.6.18

NOTE: For instructions on updating MariaDB or rebooting the Skylar One system, see the section on *Updating Skylar One* in the ***System Administration*** manual.

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

Clearing Cache Post-Upgrade

After upgrading to version 12.5.20, you should clear your system cache to remove cached items from Skylar One (SL1) and prevent several potential issues that can occur post-upgrade due to caching. To do so, go to Misc > Clear SL1 Cache.

Required Ports

Beginning with SL1 12.2.0, if you have a firewall between your Database Server, data engine, and Administration Portal appliances, you should open TCP port 8200 to facilitate communication between those appliances.

For a full list of ports that must be open on each Skylar One (SL1) appliance, see the section on *Required Ports for Skylar One* in the *Installation and Initial Configuration* manual.

Python 3.9 Execution Environment Support Deprecation

Users who are currently on 12.2.x releases and use Python 3.9 execution environments for Dynamic Applications and Run Book Automations are advised that the 12.3.0 release removed support for Python 3.9 and added support for Python 3.11. For more information, see the section *Important Notes on Creating ScienceLogic Libraries* in the ***ScienceLogic Libraries and Execution Environments*** manual.

Use of tmux When Using SSH

Starting with Skylar One (SL1) version 12.3.4, the tmux utility is disabled by default if you are on a non-STIG deployment and access a Skylar One (SL1) system using SSH. ***This is a change in behavior from versions 12.2.1.1 through 12.3.3, where the tmux utility was enabled by default.***

If you are on a STIG-compliant Skylar One (SL1) deployment, the tmux utility is enabled by default. ScienceLogic encourages non-STIG users to enable the tmux utility as well.

The utility, which is a terminal multiplexer that enables a number of terminals to be created, accessed, and controlled from a single screen, strengthens session-control mechanisms and aligns with industry-wide security practices.

If tmux is enabled, sessions are automatically locked after 15 minutes of idleness or if an unclean SSH disconnect or dropped SSH connection occurs. Upon login, Skylar One (SL1) checks for and attaches any detached tmux session if it finds them; otherwise, it starts a new session.

The utility also facilitates advanced features like scroll-back buffering with search, built-in clipboarding, multiple sessions and panes, detaching or attaching sessions, and session supervision or sharing.

To enable the tmux utility in non-STIG deployments:

1. Either go to the console of the Skylar One (SL1) appliance or use SSH to access the appliance.
2. Open a shell session on the server.
3. Type the following at the command line to edit the silo.conf file:

```
sudo visilo
```

4. Change the following line in the `[OS_HARDENING]` section of the file to enable tmux:

```
TMUX = true
```

> **NOTE:** If the `[OS_HARDENING]` heading does not already exist in the silo.conf file, you must add that immediately above the `TMUX = true` setting.

5. Save and quit the file. (`:wq`).
6. Log out of Skylar One (SL1) and then log back in. The tmux utility is now enabled.

For more information about tmux shortcuts and usage, see https://tmuxcheatsheet.com/.

Changes to High Availability and Disaster Recovery Configurations

In Skylar One (SL1) 12.3.2, the web server configuration for the vault service was updated to improve behavior between multiple database and data engine appliances. This was done to prevent false system events indicating that passive databases could not connect to Skylar One Collectors in high-availability (HA), disaster recovery (DR), or HA+DR configurations. With this update:

- On new installations, you must run the following command after all database and data engine appliances are licensed, to populate the list of allowed locations where vault could be running: `sudo /opt/em7/share/scripts/vault_add_servers.sh`. If the list is populated successfully, the output will tell you to restart nginx for it to take effect. Upon updating, this is generated automatically.
- If you are upgrading and you have not modified the default configuration file, then it will be updated to this new configuration automatically.
- If you are upgrading and you have modified the default configuration file, the updated web server configuration will be installed as **/etc/nginx/conf.d/vault.conf.rpmnew**, and you will need to merge your modifications into the new configuration and then remove the .rpmnew extension.

System Update Notes

- ***Skylar One updates overwrite changes to the configuration file /opt/em7/nextui/nextui.env***. (For more details, see 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.
- ScienceLogic recommends that you run backups of your Skylar One system before performing a system update.
- The Skylar One user interface will be unavailable intermittently during system update.
- 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 Skylar One releases during a maintenance window.
- The Skylar One 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.
- When upgrading a large number of Skylar One appliances, you might encounter an issue where the deployment summary shows that deployment timed out for many of the appliances but, upon further inspection, you discover that the appliances actually deployed correctly. This is due to a lag in the deployment status reaching the Database Server after the default timeout value of 3600 seconds (1 hour). If you check back later, the issue should fix itself. If you would rather work around this issue, you can increase the timeout value. For instructions, see the section on *Adjusting the Timeout for Slow Connections* in the "*Updating Skylar One*" chapter of the ***System Administration*** manual.
- When upgrading Skylar One on AWS stacks, you might receive an error message that the Data Engines failed to patch correctly. If this occurs, re-run the pre-upgrade tests and then run the patch again; this should result in the Data Engines updating correctly and the correct version then being reflected on the **Appliance Manager** page (System > Settings > Appliances).
- After upgrading, to ensure proper data collection, you should go to the **Appliance Manager** page (System > Settings > Appliances), locate one of the Data Collector or Message Collector appliances, and click the lightning bolt icon to force configuration push for that appliance.

Known Issues for Skylar One 12.5.20

NOTE: ScienceLogic strongly recommends that you review all Known Issues for Skylar One. For more information, see https://support.sciencelogic.com/s/known-issues#sort=relevancy.

The following known issues exist for Skylar One 12.5.20:

Deployment and Configuration

- Deployment might fail on remote appliances with a "Vault configuration error: decrypt" message. This typically occurs during the payload upload phase. To work around this issue, rerun the deployment task. (Jira ID: EM-79783)

System Upgrade

- 12.2.x and 12.3.x STIG-compliant users should contact their ScienceLogic account managers for information about upgrading to this release.
- The `Patch hook` task might fail on secondary Database Server and Administration Portal appliances when concurrent import is enabled during an upgrade. To work around this issue, manually upgrade the siloupdate RPM on the affected appliances. (Jira ID: EM-79882)
- In some cases, the secondary database appliance might become unavailable immediately after deployment completion, causing it to skip the version update. This can lead to subsequent synchronization and stability issues within the cluster. (Jira ID: EM-79857)
- When upgrading Skylar One to version 12.5.20 from 12.3.x or earlier, PowerPacks that use legacy encryption will become read-only. This issue does not impact upgrades from previous 12.5.x releases. For more information about this issue, including resolution steps, see https://support.sciencelogic.com/s/article/20806.(Jira ID: EM-73018)
- In systems that have consumed a large number of Skylar One (SL1) patch imports, the `master_filestore.storage_system_patch` database table might grow too large in size. If this occurs, then when you attempt to log in to Skylar One, you will be unable to do so and will instead receive an error message stating "The table 'organizations_log' is full" if logging in via the default user interface (AP2), or without an error message if logging in via the classic user interface. To address this issue, you should clean up any previous patch import files after deploying a new version on your Skylar One stack. For more information about this issue, see https://support.sciencelogic.com/s/article/18285. (Jira ID: EM-76040)
- System updates sometimes fail on Data Collectors due to siloupdate not updating on them. For more information about this issue, including workaround procedures, see https://support.sciencelogic.com/s/article/18508. (Case: 00543539) (Jira ID: EM-76307)
- After updating your system, you might experience an issue where the user interface is not displaying or working as intended. If this occurs, clear your browser cache. (Jira ID: EM-73859)

- During deployment, the `siloupdate-deploy-local` service might display the following log messages, although in both cases, the deployment task still ultimately completes:
    - On primary databases, it displays "Failed to enqueue storage object. Error: 'charmap' codec can't encode…" in `journalctl`. (Jira ID: EM-78481)
    - On secondary databases, it displays "ERROR 2002 (HY000): Can't connect to MySQL server on '127.0.0.1' (115)" in `journalctl` from the `sl-otelcol` post-install scriptlet. (Jira ID: EM-78479)
- An intermittent issue sometimes causes the database connection to Amazon RDS instances to become unavailable for a brief amount of time during the upgrade process, which causes deployment to be marked as failed in the user interface. If this occurs, re-run the upgrade; doing so should update all backend metadata to register as successfully completed and update the latest version in the user interface. (Jira ID: EM-66627)
- If MariaDB has an interruption during the postupdate process, the deployment is initially marked as failed. After MariaDB becomes available again, the Version Update will re-run and the appliance deployment will be marked as complete. (Jira ID: EM-78499)
- The MariaDB upgrade process might fail with an error stating that **/opt/em7/services/config.env** does not exist. This prevents the `siloupdate` utility from stopping ScienceLogic services before proceeding with the database upgrade. (Jira ID: EM-75840)
- During a patch staging retry, an appliance that completes staging successfully might still be incorrectly reported as "Failed" in the staging summary. This is a reporting error and does not necessarily mean the files were not staged. (Jira ID: EM-79600)

Access Control

- You might be able to delete an account policy even if it is currently assigned to active accounts or authentication resources, which can cause you to experience orphaned configurations and authentication issues. (Jira ID: EM-79801)
- The system erroneously enables you to create an empty access key if you click Save without filling in any fields, which can lead to invalid or non-functional permission configurations. (Jira ID: EM-76409)

Agent

- Clicking the **[Upgrade]** button on the **Agents** page results in a message indicating that the upgrade was successful immediately appearing at the bottom of the page when, in actuality, the upgrade was merely initiated. The actual upgrade process can take several minutes to complete. (Jira ID: EM-74365)
- When attempting to install a Gen 1 agent on an AWS stack, the installation user interface will display an incorrect IP for download, which will result in the download attempt failing. (Jira ID: EM-74307)
- You might experience a scenario where an agent's polled data configuration is cleared unintentionally. (Jira ID: EM-74364)
- Discovery via Streamer Prime fails for Gen 1 Agent devices if the existing device record has a class_ type of 0. This typically occurs when a device aligns before its backing PowerPack is fully installed or populated. (Jira ID: EM-79395)

- The Dead Letter Queue (DLQ) for `streamer_push` might erroneously indicate that it has exceeded its threshold. (Jira ID: EM-77934)

API

- Deleting a user account via the REST API (`DELETE /api/account/<uid>`) does not archive the account in the `master_access.accounts_deleted` database table, thus allowing the next account created to reuse the deleted account's user ID (uid). (Jira ID: SLS-2044)
- The account API prevents the `Aligned_organizations` field from being set to null. As a result, users cannot disable additional organizations for a user via the API. (Jira ID: SLS-2081)
- Attempts to update device filesystem or Dynamic Application thresholds using the API (POST/PUT) might result in a "System internal error". (Jira ID: EM-76562)

Asset Management

- When creating or editing an asset, the values entered in the ***Administrator*** and ***Technician*** fields might not be saved. (Jira IDs: EM-74258, EM-76664)

Authentication

- A known issue with session cache management might cause Skylar One to log you out unexpectedly, or prevent you from logging in again after a recent session. If you experience either issue, you can work around it by clearing the cache of your web browser before you log into the Skylar One user interface. For more information, see https://support.sciencelogic.com/s/article/13701. (Jira ID: SLUI-21011)
- The ***Single Instance Login*** setting, which can be set on the **Behavior Settings** page (System > Settings > Behavior), is not working as designed for ASCII ADFS user accounts. (Jira ID: SLS-1559)
- If you create an empty authentication resource with no values filled in and assign it to an empty authentication profile, upon logging out of the system, the user interface becomes inaccessible. (Jira ID: EM-76608)
- You might intermittently encounter an error message when attempting to change your password or submit preference changes on the **Account Preferences** page (Preferences > Account > Preferences). (Jira ID: SLS-1949)
- If you attempt to save a single sign-on authentication resource without providing the required IdP Certificate, the form displays an error and resets all other fields, forcing you to re-enter all configuration data. (Jira ID: EM-79836)

Business Services

- The **[Anomalies]** tab on the **Service Investigator** page for device services might incorrectly display devices that have anomaly detection disabled, rather than showing only those devices with anomaly detection enabled. (Jira ID: EM-62884)

- For services that have their ***RCA Options*** field enabled and have had a child service removed, Skylar One will not compute the health, availability, and risk values until the Service Topology Engine returns an updated topology, which occurs every 5 minutes by default. (Jira ID: SLUI-18853)

> IMPORTANT: Before deleting child services in a 3-tier hierarchy, check if the parent service has the ***RCA Options*** field *Enabled*, then set this field to *Disabled* if it is not already.

Credential Management and Discovery

- For an unguided device discovery, the ***Search*** box that displays for creating a new credential does not work. (Jira ID: SLUI-20777)
- When using the SNMP Public V2 credential to discover devices, you might see an unhandled exception in the system log near the end of the discovery session, despite the devices being discovered successfully. (Jira ID: EM-59380)
- When selecting two or more SNMP credentials to discover a device, if the first credential with the lower ID number contains incorrect information and the second credential with a higher ID number contains the correct information, the discovery logs will not be able to get an SNMP response. (Cases: 00289639, 00292649, 00422558) (Jira ID: EM-39681)
- The **Credentials** page in the default user interface (AP2) fails to display credentials that are not aligned with an organization, but displays these credentials correctly in the classic Skylar One user interface on the **Credential Management** page. (Jira ID: SLUI-20947)
- On the **Credentials** page, if you have more than 50 credentials and at least one of the first 50 credentials is not aligned with an organization, the page will display duplicates of these credentials. (Jira ID: SLUI-20947)
- You might get an error when trying to open SOAP/XML credentials that have been imported from a PowerPack. (Jira ID: EM-74291)

Dashboards

- In dashboards created in the default user interface (AP2), Devices and File Systems widgets might display incorrect data due to a known mathematic scaling issue. To work around this issue, consult the same metrics in the **[Performance Metrics]** tab of the **Device Investigator**, or in the **[Performance]** tab of the **Device Reports Panel** in the classic user interface. (Jira ID: EM-79597)
- Creating an **Interface** widget with the *Leaderboard* visualization, applying an advanced filter, and adjusting the data time span using the ***Time span filter*** results in an error. (Jira ID: SLUI-22200)
- When editing the scale prefix of a **Device** widget using the *Leaderboard* visualization, the ***Storage Used*** column does that update in that widget's table. (Jira ID: SLUI-22198)
- When editing an **Events** widget and setting the ***Refresh Mode*** field to *None*, the widget's events table shows the refresh mode as automatic, despite the change. (Jira ID: SLUI-21947)
- In classic dashboards, if you create a Traffic Light widget, the ability to control context in other subscribing widgets is not working as intended. (Jira ID: EM-76527)

- You might encounter incorrect values in Business Service dashboards and device metric trend charts when viewing physical memory aggregations due to a mapping error in the normalized database views causing physical memory sum data to track CPU activity instead. (Jira ID: EM-79948)
- The "Shared: Interface Billing" dashboard might generate a large number of console errors when loading, particularly on STIG systems. (Jira ID: EM-77611)
- The "Forecast: Contextual Device" widget on classic dashboards might fail to display data and instead show a perpetual loading state, preventing you from viewing CPU forecast graphs for selected devices. (Jira ID: EM-80162)
- When imported via a PowerPack, the "Custom Table: Journal Apps" dashboard widget's settings might appear, but selections do not display correctly. (Jira ID: EM-76599)
- The "Event Counts by Severity" widget fails to display any data when the event state is set to "Acknowledged," even if acknowledged events exist in the system. The widget correctly displays data only when set to "Unacknowledged." (Jira ID: EM-76286)

Data Collection and Retention

- If you include an invalid or incorrectly typed value in silo.conf, such as a word or phrase where it should be an integer, the data pull process crashes and cannot start. (Jira ID: EM-74238)
- When aligning a Dynamic Application that discovers a dynamic component map tree using Latin-1 and UTF-8 encoded device names and identifiers, you might receive "Illegal mix of collations" data storage errors in the system logs. (Jira ID: EM-74263)
- Some performance metrics and alert formulas display values with excessive decimal places instead of rounding to two. (Jira ID: EM-79390)
- Disabling interface collection can cause the "Interface Magic" service in the collector pipeline to crash with a ZeroDivisionError. (Jira ID: EM-72808)

Device Management

- On the **Devices** page, when sorting your search by the ***Organization*** column, the inventory table will sort by ***Organization ID*** instead. (Jira ID: SLUI-21459)
- The assigned organization for devices might not always update, even after performing a bulk alignment organization action on the **Devices** page. To work around this issue, refresh your browser immediately after completing the bulk alignment action. (Jira ID: SLUI-21483)
- The column widths on the **Device Investigator** page do not adjust when resized. (Jira ID: SLUI-20081)
- Filtering the ***Collector Groups*** column on the **Device Investigator** page with multiple group names can cause the page to not load correctly. (Jira ID: SLUI-21035)
- When sorting by columns on the **Device Investigator** page in Firefox, the table might continuously attempt to retrieve results unsuccessfully. (Jira ID: SLUI-21095)
- The **Device Categories** page (Devices > Device Categories) fails to load properly whenever there is a category with a null ID. To work around this issue, go to the **Device Categories** page (System > Customize > Device Categories), locate the category with the null ID, and then remove that category by clicking the delete icon () next to the category.(Jira ID: SLUI-20731)

- The number of unacknowledged events in the ***Device Overview*** panel of the **Device Investigator** page does not update despite acknowledging alerts on a device. To work around this issue, add a new "unackEvents" subquery to the "Device Insights" query, then use that subquery to collect and retrieve information on unacknowledged events. (Case: 00471966) (Jira ID: SLUI-20858)
- When attempting to bulk delete devices or device components, a dialog message might indicate that some or all of the devices or components failed to delete, when in fact they were actually deleted. (Jira ID: EM-74351)
- Device reports generated from the **Classic Devices** page (Devices > Classic Devices) might result in an "Uncaught TypeError" message appearing in the browser console, despite most reports generating correctly. (Jira ID: EM-74225)
- HTML device journal reports might generate with no data. (Jira ID: EM-74286)
- The checkbox to select all items on the **Device Dashboards** page (System > Customize > Device Dashboards) might select only the items that appear on your current page rather than all available device dashboards. (Jira ID: EM-74215)
- Event messages in device logs for multi-match events might display with the incorrect event ID. (Jira ID: EM-74409)
- After upgrading to 12.5.20, you might see multiple **[Attributes]** tabs on the **Device Administration** panel in the classic user interface. This issue does not affect ISO deployments. (Jira ID: EM-79933)
- The trend checkboxes on the **[Performance]** tab of the **Device Reports** panel do not always accurately update the graph display. Additionally, unchecking a trend does not hide the area fill, and re-checking it causes the trend to disappear entirely. (Jira ID: EM-79776)
- After successfully updating a device class on the **Device Properties** page, the system might erroneously display a "Device Class/Type Update Failed" error message. (Jira ID: EM-79472)
- After un-aligning a custom attribute from a device, the Custom Attribute Subscriber modal might incorrectly display all devices in the system. This is a visual issue only; it does not impact the actual alignment status. (Jira ID: EM-76508)
- Saving changes to a device's custom attributes might cause unset integer-based attributes to be automatically set to 0 rather than remaining null. (Jira ID: EM-76481)
- Changes made to a device template during the "Modify By Template" process are not saved to the template itself, even when you select the ***Save When Applied & Confirmed*** checkbox, preventing you from updating persistent templates while applying them to devices. (Jira ID: EM-74696)
- Nightly discovery might trigger a "TypeError" exception when updating interfaces for PowerShell-monitored devices. This issue appears in the storage process logs and affects the accuracy of interface IP updates. (Jira ID: EM-73781)
- You cannot submit bulk administrative requests on the **Device Components** or **Device Hardware** pages in the classic user interface. A “Loading…“ message appears on the page and the submitted request fails. For more information, see https://support.sciencelogic.com/s/article/18343. (Jira ID: EM-73389)
- The "System Vitals Performance" graph might display a blank screen if a device is missing a collection label. (Jira ID: EM-77642)

Enterprise Key Management Service (EKMS)

- You might experience an issue where the EKMS state and vault contents disappear from your Data Collectors. If this occurs, you can work around the issue by going to the **Appliance Manager** page (System > Settings > Appliances) and manually running the ***Enterprise Database: Collector Config Push*** process to run on the impacted appliances. (Jira ID: EM-78522)
- Configuration backups do not include the `master.vault_token` table, so restoring from a config backup will not restore Vault token data. Users who depend on full database restores for recovery should ensure this information is backed up by other means. (Jira ID: SLS-1948)

Events and Alerts

- Deviation alerting does not support the use of double quotes in indices. To work around this issue, use single quotes. (Jira ID: EM-72050)
- The events on the **Events** page cannot be sorted by the ***Organization*** column. (Jira ID: SLUI-20903)
- From the **Event Policies** page (Events > Event Policies), you can delete only a single event policy at a time, even if you select multiple event policies for bulk deletion. (Jira ID: SLUI-20853)
- Event messages derived from incoming email substitute "\xc2" for the character Â. (Jira ID: EM-73551)
- You might occasionally be prevented from taking certain actions on the **Classic Events** page such as clearing or acknowledging events. (Jira ID: EM-79905)
- On the **Event Statistics** page, after changing the graph type (for example, from line to stepline), the checkboxes for showing or hiding trends and selections might stop functioning. The boxes remain unresponsive even if the graph type is changed back to the original setting. (Jira ID: EM-79888)

Global Manager

- The devices on the **Devices** page in Global Manager systems cannot be sorted by the ***IP Address*** column. (Jira ID: SLUI-21108)
- On Global Manager systems, the *View Event Policy* option in the **Actions** menu (⋮) on the **Events** page does not work as expected. (Jira ID: SLUI-21133)
- On Global Manager systems, the **Events** page does not display events from child stacks. To work around this issue, clear all system caches on both the child stacks and the Global Manager parent stack, then restart the NextUI service. (Jira ID: SLUI-21134)

GraphQL

- The "harProviderOnDemanProcessing" GQL query incorrectly creates a service table in the "data_har" database when executed with invalid or non-existent service IDs. (Jira ID: SLUI-21135)
- Clicking the **[Run Now]** button for any Dynamic Application on the **[Collections]** tab of the **Device Investigator** will display the following GQL error message in the Skylar One server console: "Variable "$procId" of non-null type "ID!" must not be null." (Jira ID: SLUI-21070)

Logging

- System logs related to backup start, stop, and failure events are not being generated as intended. While a backup might complete successfully, you cannot use the system logs to validate backup timelines or troubleshoot issues. (Jira ID: EM-79904)
- A known issue might cause several log configuration files to conflict, which could cause you to see errors for the `sl_vault` and `slsctl` logs or potentially block log rotation in some cases, depending on the order in which the files are executed. To work around this issue, delete the config files `~sl_vault` and `~slsctl`. (Jira IDs: SLS-1105, EM-62134)
- In debug mode, Config Push logs might show repeated "OperationalError: Partition management on a not partitioned table is not possible" messages. This log does not affect system functionality or appear in the user interface. (Jira ID: EM-79813)
- The `php-error.log` does not track actions for non-administrator local users. (Jira ID: EM-77688)

Monitoring Policies

- New Domain Name monitoring policies are successfully created only if you enter all required fields. If a required field is missed and you attempt to save, you must exit out of the modal and start again. (Jira ID: EM-80340)

Platform

- You might receive an "Internal Server Error" when attempting to configure an IP address on a heartbeat interface using the Web Configuration Tool. To work around this issue, configure heartbeat IP addresses using the command line interface. (Jira ID: EM-79886)

PowerPacks

- When upgrading Skylar One to version 12.5.20 from 12.3.x or earlier, PowerPacks that use legacy encryption will become read-only. This issue does not impact upgrades from previous 12.5.x releases. For more information about this issue, including resolution steps, see https://support.sciencelogic.com/s/article/20806.(Jira ID: EM-73018)
- In STIG deployments, you might not be able to add AP2 content objects to PowerPacks. (Jira ID: EM-78797)
- Due to a known issue, you might need to manually upgrade to the following PowerPack versions after installing or upgrading to Skylar One 12.5.1 or later:
    - "Net-SNMP Base Pack" PowerPack v103 (Jira ID: EM-73518)
    - "Microsoft Base Pack" v110 (Jira ID: EM-73516)
    - Microsoft: Windows Server" v118 (Jira ID: EM-73516)
- When installing or importing a PowerPack, you might not be able to adjust the PowerPack's embedded license or license key type. (Jira IDs: EM-71507, EM-72515, EM-72716)
- If a PowerPack installation fails, you are now prompted to contact ScienceLogic Support rather than being stuck on the installation page. (Jira ID: EM-77519)

- You might encounter an issue that prevents you from deleting a PowerPack through the user interface. (Jira ID: 76085)
- After updating the "ScienceLogic: Default Internal Events" PowerPack, the PowerPack Manager page might incorrectly display version 12.5.2 instead of version 100. (Jira ID: EM-79963)
- If you use the "IBM: DB2" PowerPack, component information might not be available when the "IBM DB2: Dynamic Application Alignment" Run Book Policy and Action is run and the name of the parent component is unavailable to the Run Book Action. This will result in the additional Database-level Dynamic Applications being aligned with the Default SNMP credential. For more information, see https://support.sciencelogic.com/s/article/20838.
- When creating a new PowerPack, attempting to save modifications without first refreshing the page might trigger an "Error: License key is not valid" message. To work around this issue, refresh the PowerPack Editor page. (Jira ID: EM-73201)

Reporting

- If you create a PDF report about a single device from the **Devices** or **Classic Devices** pages (Devices > Classic Devices, or Registry > Devices > Device Manager in the classic user interface), some tables might not display as intended. (Jira ID: EM-79587)
- The DEV_REPORT_CREATOR access hook currently only restricts the print icon on the classic **Device Registry** page. Users without this hook can still generate reports on the **Device Properties** page, on the **Device Summary** page, and in the default user interface (AP2). (Jira ID: EM-79956)
- If you attempt to load a performance report with images, it will fail if the *Percentile* option is enabled on the graph. Reports without images are unaffected, but any report including visual graph data will fail to load in this state. (Jira ID: EM-79893)
- The ***Notes*** section of generated User Report PDFs might display raw HTML tags instead of the intended rich text formatting. (Jira ID: EM-79867)
- Generating Asset List Unsafe (v1.62) quick reports might fail for multiple formats, including PDF, ODS, and XLSX. The report status is marked as "Failed" in both the default and classic user interfaces. (Jira ID: EM-79825)
- On the **Report Output Styles** page (Reports > Management > Report Output Styles), the ***Style Name*** column has values with non-breaking spaces that are displayed as HTML code. This issue does not prevent you from creating or running reports. (Jira ID: EM-80296)

Schedule Management

- A large number of stored procedure calls over time can lead to memory growth in the scheduler process as well as delays in processing schedules. (Jira IDs: EM-76720, EM-75045)

Skylar AI

- The `sl-otelcol-mgmt.py` script shows blank values for the `sl-otelcol`, `em7-platform-core`, and `nextui` versions when connecting to Skylar AI, making it unclear which versions are actually installed, even though the status output reports the versions correctly. (Jira ID: EM-78655)
- Anomaly Detection and Predictive Alerting are not created for Dynamic Applications that are not associated with a PowerPack. (Jira ID: EM-73074)

- Skylar AI Connection debug always runs on the node the connection was initially configured with, even after a failover. When this occurs, it causes a false negative status. (Jira ID: EM-78209)
- Supplying an open telemetry collector endpoint URL with a trailing slash can cause the endpoint to be rejected or fail validation. As a workaround, remove any trailing slash from the URL before saving. (Jira ID: EM-77859)
- If the Skylar AI endpoint is unavailable, the Open Telemetry Collector might receive data from Skylar One but fail to send it to Skylar AI without notifying the sender. (Jira ID: EM-78343)

Subscription Billing

- The Bandwidth Billing interface filter might become unusable after the initial search, even if the filter is cleared. To work around this issue, click Reset. (Jira ID: EM-76511)

System Administration

- When editing a collector group from the **Collector Groups** page (Manage > Collector Groups) by clicking its **Actions** menu (ellipsis icon) and unchecking two or more organizations in the ***Limit access to specific organizations*** field of the **Edit Collector Group** modal, the ***Organizations*** column on the **Collector Groups** page will show that only one organization was deselected, even if multiple were. (Jira ID: SLUI-22167)
- Message Collectors on the **Collector Groups** page (Manage > Collector Groups) cannot be sorted by the ***Message Collectors*** column. (Jira ID: SLUI-22099)
- Default account policies might automatically align with the organization that is alphabetically ranked lowest, potentially causing unexpected policy assignments. (Jira ID: SLS-2084)

User Interface

- In the Skylar One user interface, the End User License Agreement (EULA) page is displayed on all pages that were iframed from the classic user interface, even after the user agrees to the EULA. This issue is occurring for ADFS, CAC, and AD authentication methods. (Jira ID: EM-67851)
- The **[Expand]** and **[Contract]** buttons are not working as intended on the **Dynamic Application Collections** page (Devices > Device Manager > wrench icon > Collections). You can still expand and contract individual items on the page. (Jira ID: EM-64420)
- The **Access Keys** page (System > Manage > Access Keys) might not count administrator users in the value displayed in the ***# Aligned Users*** column. To work around this issue, go to the **Account Permissions** page (Registry > Accounts > User Accounts > wrench icon) for the administrator users and re-save their permissions. (Jira ID: EM-74241)
- When you bulk-select multiple event policies to align with a run book automation policy, additional event policies that you did not select might become aligned with that automation policy as well. (Jira ID: EM-70690)
- On the **Custom Attributes** page (Manage > Custom Attributes), you might not be able to view more than the first 20 custom attributes unless you zoom in or change the size of your browser to force Skylar One to fetch additional attributes. You also might not be able to select the "Select All" checkbox on the page. (Jira IDs: SLUI-21449, EM-74251)
- Name changes to nodes on the **Nodes** page (Manage > Nodes) are not saved. (Jira ID: SLUI-22248)

- In the default user interface (AP2), when opening the **Account Permissions** page (Registry > Accounts > User Accounts > wrench icon) for an existing user account, the ***Theme/Brand*** drop-down field does not initially display on the page. To work around this issue, refresh the page. (Jira ID: EM-76478)
- On the **OID Browser** page (System > Tools > OID Browser), the *Where Symbolic is like* drop-down option for the ***Search where*** field might not work as intended. (Jira ID: EM-74326)
- In the classic user interface, the filter for ***Edit Date*** is not working as intended on the **Inbound Email** page (Registry > Events > Inbound Email). (Jira ID: EM-75291)
- When using a dark mode theme, if you click on a calendar icon to select a date, some dates might not be visible on the calendar. (Jira ID: EM-76629)
- The final row might not appear on the **Select Objects** page (System > Customize > Select Objects) when viewed in the default user interface (AP2) using a Firefox browser. (Jira ID: EM-74222)
- When using Active Directory Federation Services (ADFS) to authenticate, the system opens to the default landing page rather than the last page the user visited. (Jira IDs: SLS-1764, SLS-1765)
- If you create a new network on the **IPv4 Networks** page (Registry > Networks > IPv4 Networks), it might not appear on the page after you save it. However, the network should still be available for selection in other operations. (Jira ID: EM-76666)
- In AWS deployments, the ***% Used*** column on the **IPv4 Networks** page (Registry > Networks > IPv4 Networks) might erroneously display values over 100%. (Jira ID: EM-76672)
- When saving a change on the **Quality of Service Threshold Defaults** page (System > Settings > Thresholds > Quality of Service), the page might continuously state that it is saving until you refresh the page. Despite this messaging issue, the system should save the changes correctly. (Jira ID: EM-74336)
- If you select a normalized option for a service usage graph, the **Date Range Selection** pane changes to the year 1970 and the graph will not be visible, even if you correct the start and end dates. (Jira ID: EM-74331)
- When creating an uptime OID, the system might display an error alert if the OID is invalid but still save the OID anyway. (Jira ID: EM-74227)
- On the **Process Manager** page (System > Settings > Admin Processes), if you deselect the appliance type(s) that a process should run on and then save that change, it will prevent that process from starting in the future. (Jira ID: EM-74236)
- If you click the **[Save]** button on the **User Policy Properties Editor** page (Registry > Accounts > User Policies > create or edit) for an existing user policy, the system mistakenly saves the user policy again as a separate entry rather than overwriting the existing policy with the updated information. (Jira ID: SLS-1773)
- The ***Sharing Permissions*** drop-down field on the **Appliance Settings** page incorrectly displays *Private* by default, even though the system treats the setting as *Shared*. (Jira ID: SLS-1870)
- Going to some pages or performing certain actions in the classic user interface might trigger "PHP Deprecated" notices in the error logs. These notices do not impact page functionality and can be safely ignored. (Jira IDs: EM-79898, EM-75356, EM-75310)
- The **[Save As]** button for Database type credentials might be invisible or partially obscured when using the Firefox browser. This issue specifically affects systems that have upgraded through multiple versions and is not present in other browsers. (Jira ID: EM-79885)

- After performing an initial search in the **Bandwidth Billing Policy Editor**, the **[Filter]** button can become inactive, preventing users from changing the filter text. To work around this issue, click the **[Reset]** button. (Jira ID: EM-79847)
- Sorting by the ***Organization*** column on the **SSL Certificate Monitoring** page might produce inaccurate results due to the system sorting by the organization ID rather than the organizations name. (Jira ID: EM-79843)
- In classic Service Provider Utilities, Service Usage policies might fail to save selections if you choose more than one device. (Jira ID: EM-79809)
- On the **User Policy Membership** page (Registry > Accounts > User Policies) in the classic user interface, sorting by the First Name, Phone, or Mobile columns might cause all rows to disappear, preventing administrators from managing policy members while the sort is active. (Jira ID: EM-74702)

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

International: +1-703-354-1010