



STORMSHIELD



STORMSHIELD LOG SUPERVISOR

RELEASE NOTES

Version 1

Document last updated: July 13, 2023

Reference: [sls-en-release_notes-v1.1.1](#)



Table of contents

Change log	3
Compatibility	4
New features and enhancements in version 1.1.1	5
Version 1.1.1 bug fixes	6
Plugins and Vendors	7
Known Issues	9
Documentation resources	10
Installing this version	11
Downloading the SLS patch	11
Updating SLS from version 1.1.0 to version 1.1.1	11
Updating SLS from version 1.0.0 to version 1.1.1	12
Previous versions of SLS 1	13
Contact	30

In the documentation, Stormshield Log Supervisor is referred to in its short form: SLS.
This document is not exhaustive and minor changes may have been included in this version.



Change log

Date	Description
July 13, 2023	- Added SES Evolution compatibility - Section Known Issues added
May 24, 2022	- Added Microsoft Hyper-V Server 2022 compatibility
November 25, 2021	- SLS 1.1 Release - Added VMWare ESXi server 7.0 compatibility - End of support for VMWare ESXi server 6.5



Compatibility

The following components are compatible with SLS version 1.1.1.

Stormshield products

Stormshield Network Security	Version 3.7 and higher
Stormshield Endpoint Security	Version 2.4.3 and higher

Hypervisors

VMware ESXi	Versions 6.7 and 7.0
Microsoft Hyper-V Server	Versions 2016, 2019 and 2022

Filesystem

ext4

Web browsers

In order for the firewall administration interface to operate optimally, you are advised to use the latest versions of Safari, Google Chrome and Mozilla Firefox (ESR version - Extended Support Release). For further information on these versions, please refer to the relevant vendors for the life cycles of their products.



New features and enhancements in version 1.1.1

SNS Stormshield application updates

- Stormshield application has been updated. You must now use the *SNS Stormshield* application to normalize Stormshield events and analyze the data using pre-set dashboard views.
- *SNS Stormshield* v5.1.3 application is bundled with this patch.
- The *sent* and *rcvd* fields are now renamed:
 - *sent* is now *sent_datasize*,
 - *rcvd* is now *received_datasize*.



Version 1.1.1 bug fixes

SNS Stormshield application

- Users could not uninstall *Stormshield* application. Updating SLS to version 1.1.1 automatically removes the previous *Stormshield* application.
- If you have updated SLS to version 1.1.1 and were using the previous SNS built-in dashboards without any customization, you must manually remove the old dashboards and add the new ones. For more information, refer to the [Updating SLS from version 1.1.0 to version 1.1.1](#) section.



Plugins and Vendors

Plugins

Plugins allow to use specific applications. The following applications have been bundled with Stormshield Log Supervisor 1.1.1.

Name	Version
Ascii Converter	3.0.0
Base 16	3.5.0
Clean Char	3.1.0
Codec	3.1.0
CSV Enrichment Source	5.1.0
DNS	3.1.0
GeoIP	5.0.2
Grok	3.0.1
InRange	3.1.0
IPtoHost Enrichment Source	5.0.0
LDAP Enrichment Source	5.0.0
LogPoint	5.2.2
MacToVendor	5.0.0
Mitre Dataset Updater	5.0.0
Regex	5.0.0
SNS Stormshield	5.1.3
Syslog Forwarder File Fetcher	5.0.0
Threat Intelligence	5.2.0

Vendors

Vendor Alert Rules

1. Threat Intel Internal Machine Connecting to Multiple IOCs
2. Threat Intel Excessive Denied Connections Attempt from IOC
3. Threat Intel Connections with Suspicious Domains
4. Threat Intel Allowed Connections from Suspicious Sources
5. Threat Intel IOC Connecting to Multiple Internal Machines
6. License Expiry Status



7. License Invalid
8. License Grace State

Vendor Dashboards

1. SNS Web
2. SNS Threats
3. SNS System
4. SNS Overview
5. SNS Interface
6. SNS Hardware
7. SNS Data
8. SNS Applications

Vendor Report Templates

1. SNS Web and Applications Activity
2. SNS Threats Activity
3. SNS System Overview
4. SNS General Activity

Vendor Search Templates

1. SNS Monitor



Known Issues

The up-to-date list of the known issues related to this SLS version is available on the Stormshield [Knowledge base](#). To connect to the Knowledge base, use your [MyStormshield](#) customer area identifiers.



Documentation resources

The technical documentation resources are available in the documentation base on the [Stormshield technical documentation](#) website. We suggest that you rely on these resources for a better application of all features in this version.



Installing this version

Downloading the SLS patch

1. Log in to your [MyStormshield](#) personal area.
2. Go to **Downloads > Downloads**.
3. Select **Stormshield Log Supervisor** from the suggested categories.
4. Locate the SLS patch (.pak file) you want to download and click on its name to download it.
5. You can check the integrity of the downloaded file using one of the following commands:
 - Linux operating system: `sha256sum <filename>`
 - Windows operating system: `CertUtil -hashfile <filename> SHA256`

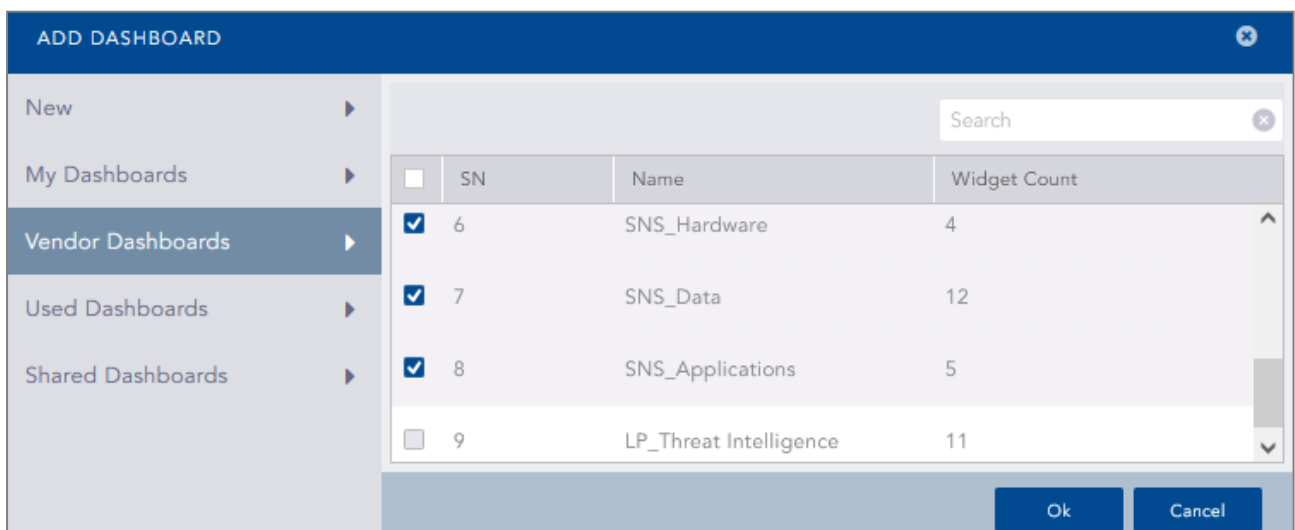
Next, compare the result with the hash indicated in the MyStormshield interface. To view it, click on **Display** in the **SHA256** column of the file in question.

Updating SLS from version 1.1.0 to version 1.1.1

1. Download the SLS version 1.1.1 patch.
2. Go to **Settings >> System >> Updates** in the SLS user interface.
3. Click **Upload Patch**. Browse for the downloaded .pak file and click **Upload**. Please wait a few minutes as the file may take some time to upload.
4. Click the **Install Update** icon from the **Actions** column.
5. Wait a few moments for the update to end. Note that the **Install Logs** window will not notify you when the installation is complete. Go back to the **Updates** page to check that the *Status* of the SLS patch is *installed*.
6. To access the new version of SLS, log out from the user interface and then log in.

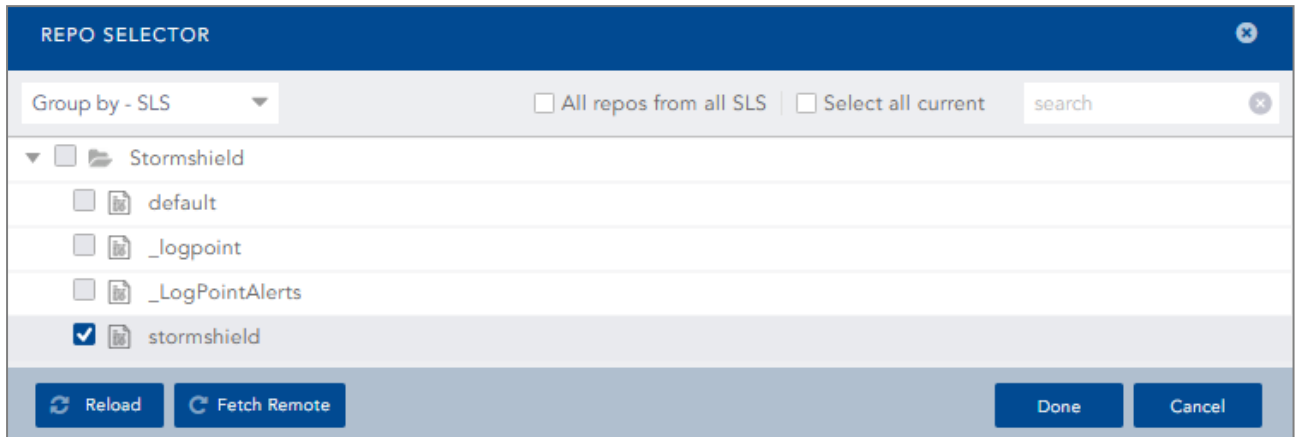
If you were using SNS built-in dashboards without any customization, you must manually remove the old dashboards and add the new ones:

1. Go to **Dashboard** from the navigation menu.
2. Remove all SNS dashboards by clicking the **Remove Dashboard (X)** icon and clicking **Yes**.
3. Click the + icon to add the new dashboards and click **Vendor Dashboards**.
4. Select all **SNS** dashboards, and only SNS dashboards, then click **Ok**.





5. Select all **SNS** dashboards again, and click **Choose Repos.**
6. Unselect **Stormshield**, select **stormshield** (in lower case), then click **Done.**



7. Click **Ok** to add the new dashboards.

Updating SLS from version 1.0.0 to version 1.1.1

To install the SLS version 1.1.1 patch, you must first update SLS to version 1.1.0.

1. Download the SLS version 1.1.0 patch.
2. Go to **Settings >> System >> Updates** in the SLS user interface.
3. Click **Upload Patch**. Browse for the downloaded *.pak* file and click **Upload**.
4. Click the **Install Update** icon from the **Actions** column.
5. Wait for the update to end. This takes several minutes. You can refresh the **Updates** page at anytime to see if the update is complete. When completed, the *Status* of the SLS patch is *installed* and the machine restarts.
6. Refresh the page to access the new version of SLS.
7. Follow the instructions given in the [Updating SLS from version 1.1.0 to version 1.1.1](#) section.



Previous versions of SLS 1

In this section, you will find the new features, resolved vulnerabilities and fixes from previous versions of SLS 1.

1.1.0	New features	Resolved vulnerabilities	Bug Fixes
1.0.0	New features		



New features and enhancements in version 1.1.0

! IMPORTANT

- You may experience a significantly increased patch upload and upgrade time because of the OS upgrade. You can check the progress in the installation logs from **Settings >> System >> Updates**.
- Each SLS instance installed from the upcoming virtual images uniquely generates an SSH key, a TLS certificate, and a TLS key on the first boot.

Report Template

You can now benefit quick report setup with built in Stormshield SNS Report Templates. Four different report templates are proposed and directly available when starting the solution:

- General Activity report – General information related to events, traffic, and users
- Threat report – Information focused on threats that have been detected and blocked by the firewalls
- Web & Applications report – Overview of web and application usage
- System report – Information focused on firewall systems and administration activities

System Upgrade

The base operating system has been upgraded to Ubuntu 20.04. Also, the Zettabyte File System (ZFS) has been upgraded to v0.8.3, and the kernel has been upgraded to v5.4.0-42. Additionally, support for TLS v1.3 has been added to Nginx as OpenSSL is upgraded after the OS upgrade.

Dashboard Sharing Mechanism with Role-Based Access

You can now share a dashboard with different users and give them the read, edit, or full permissions. A user given full permission can read, edit, delete, and share the dashboard with other users. Any changes made in the dashboard is visible to all the shared users.

Alert Sharing Mechanism with Role-Based Access

You can now share alert rules with different users and give them read, edit, or full permissions. A user given full permission can use, clone, edit, delete, and share the alert rules with other users. Incidents for each shared user and owner are triggered independently.

Role-based Access Control for Report Templates

You can now share a report template with all the users in the system and give them the read, edit, or full permissions. A user given full permission can view, edit, delete, and share the report template with other users. Any changes made in the report template are visible to all the shared users.



Role-based Access Control for Search Templates

You can now share a search template with all the users in the system and give them the read, edit, or full permissions. A user given full permission can view, edit, delete, and share the search template with other users. Any changes made in the search template are visible to all the shared users.

Manual Trigger for Alert Notifications

You can now configure the alert notifications to be sent either manually or automatically to the external MDR services of your choice. Also, you can re-send the manually triggered alert notifications for the same incident multiple times.

After upgrading, the notification triggers for all the previously configured alerts are set to **Automatic** by default. You can set them to **Manual** by going to **Knowledge Base >> Alert Rules >> Configure Notifications** and changing the trigger configuration.

Alert Rule Categorization

You can now categorize the alert rules and incidents according to the MITRE ATT&CK Framework and also provide custom metadata information while creating an alert or an incident. You can configure the following fields to categorize the alerts and incidents:

- Attack Tag
- Attack Category
- Metadata
- Log Sources

Incident Notification Configuration

You can now configure the following settings from the *Send Incident(s) for Investigation* dialog box while sending multiple incidents for investigation:

- Emails
- Subject
- Enable or Disable Search link
- Enable or Disable logo
- Add custom logo

Any pre-configured settings for email notifications are replaced with the configurations set in the dialog box for the particular instance. All the selected incidents will be grouped in a single email.

Alert Rule Views

You can now view the alert rules in two different ways: Tabular view and Coverage view. The tabular view displays additional columns to list the *Log Sources*, *Attack Category*, and *Attack Tag* associated with the alert rules on top of the default alert view and the coverage view displays the categorization of alert rules based on different attack categories and attack tags.



Other Enhancements in Alerts and Incidents

- You can now monitor the status of each incident using the *Incident ID*. The *Incident ID* is a unique ID generated when an incident is generated. You can view the *Incident ID* and corresponding *AlertRule's AlertRule ID* by clicking the *Incident Data* option in the Incident page for each incident
- You can insert Jinja syntax to customize the *Subject* field and receive dynamic subjects in the email notification.
- You can now send incidents generated from Search, Dashboard, and Search template for investigation manually.
- You can now view the incident data in a new browser tab by clicking the *Open in new tab* icon on the incident page. The icon redirects you to the search page and shows the log messages that triggered the incident. You can also view incident information like *Alert Name*, *Incident ID*, and *Incident Timestamp* from the *Incident Info* icon on the redirected tab.
- You can now view the attack tags and attack categories associated with the alert rule that generated the incidents on the incident page. You can filter these incidents according to the *Attack Category*, *Attack Tag*, and *Log Sources*.
- A validation message now appears in the *Search Interval* field when you submit a search interval value which is not a factor of the *Time-range* configured when creating an alert.

Enhancement in Visualization

The *Display chart* now supports an additional response type, *Single Aggregation with Grouping* in *Search*, *Dashboards*, and *Search Templates*.

Enhancements to the ATT&CK Chart

- Support for the MITRE ATT&CK Framework v9, which includes new techniques, sub-techniques, and technique changes that improve the effectiveness of the ATT&CK chart.
- The ATT&CK chart now displays the attack sub-techniques in addition to the attack techniques provided by [Mitre](#).
- You can now further investigate an ATT&CK chart by grouping the results based on IP, user, or workstation.
 - The chart displays the corresponding entity icon if you group the results by IP, user, or workstation. Clicking the icon displays a bar chart showing the distribution of results by the corresponding entity.
 - The chart displays an info icon on each cell if you group the results by IP, user, or workstation. Clicking the icon displays a bar chart showing the distribution of results by the corresponding entity for each technique or sub-technique.
- You can now further drill-down and filter the attack data by clicking on each technique, sub-technique, and tactic.
- The Mitre Dataset Updater v5.0.0 has been added to support additional features of the ATT&CK chart.

Other Enhancements

- You can now disable the search link present in the email notification of an alert.
- The pattern and example text boxes while adding or editing signatures in normalization packages are now re-sizeable and expand up to 200 px height.



- You can now reorder dashboards in search templates.
- You can now provide additional parameters to specify the required date-time format in the *datetime* filter used in the Jinja template while setting up alert notifications.
- You now have the option to remove the logo in the alert notifications sent via email. You can also upload a user-defined logo in JPG or JPEG format.
- Color legends are displayed on the *Incident page*, indicating the severity level of incidents as shown in the following table:

SN	Severity Level	Color
1	Critical	Red
2	High	Purple
3	Medium	Blue
4	Low	Gray

- The patches are now decrypted using the Advanced Encryption Standard (AES -256).
- A new SNMP OID has been added to monitor the number of ongoing search processes.
- Syslog collector now supports only TLS 1.2 and TLS 1.3. Additionally, SEED with 128 and 256 bit CBC ciphers have been removed. Also, various cipher suites have been updated to prevent potential security vulnerabilities.
- The *change-syslog-port* li-admin command has been introduced to toggle Syslog ports between 514 and 601.
- The message displayed for empty search results has now been made consistent throughout the application.
- The compression format of the indexes has been changed to BGZF. The system can now compress index files with sizes greater than 2GB.
- The sFlow collector now recognizes the flow samples of TCP and IPv6. Moreover, support for VLAN tags, IPv4 flags, and TCP flags has also been added.
- You can now view a detailed warning message when you submit restricted HTML tags in the input fields. You can use HTML tags to customize the structure of incident data and alert notifications. Some HTML tags are restricted for security against the Cross-site scripting security vulnerability.
- The list of processing policies now displays the normalization, enrichment, and routing policies configured for each processing policy.



Resolved vulnerabilities in version 1.1.0

The following vulnerabilities have been addressed:

- An XSS vulnerability in the python-lxml's clean module that could allow attackers to run arbitrary HTML/JS code.
- A vulnerability in urllib3 before v1.25.9 that could allow CRLF injection if attackers control the HTTP request method.
- A vulnerability in the PyYAML library that could lead to arbitrary code execution.
- A vulnerability in httplib2 before v0.19.0 that could cause a Denial of Service.
- A CRLF injection vulnerability in httplib2 that allowed %xx quote of space, CR, LF characters in the URI.
- A vulnerability in python-cryptography v3.2 that could allow Bleichenbacher timing attacks in the RSA decryption API, via timed processing of valid PKCS#1 v1.5 ciphertext.
- A vulnerability in Pillow before v8.1.0 where SGIRleDecode had a buffer over-read when decoding crafted SGI RLE image files and PcxDecode had a buffer over-read when decoding a crafted PCX file.
- Vulnerabilities in OpenLDAP before v2.4.57 led to an invalid pointer free, a double free, a memch->bv_len miscalculation, and slapd crashes that could cause a Denial of Service.
- The use-after-free vulnerabilities in the Linux kernel that could cause a Denial of Service.
- A vulnerability where the LIO SCSI target implementation in the Linux kernel could lead to the exposure or modification of sensitive information.
- A vulnerability in the NVIDIA GPU display driver for the Linux kernel that could cause a Denial of Service or escalate privileges.
- A vulnerability in the Linux kernel where memory corruption could be exploited to gain root privileges from unprivileged processes.
- A vulnerability where FasterXML Jackson Databind did not have entity expansion secured properly, that could allow XML external entity (XXE) attacks.
- The Improper Input Validation vulnerability in Apache HttpClient before v4.5.13 and v5.0.3.
- A vulnerability in OpenSSH before v8.5 that could allow attackers to unconstrained ssh-agent socket access on a legacy operating system or forward an agent to an attacker-controlled account or host.
- A vulnerability in the Linux kernel's seq file.c in the Filesystem layer that could allow an unprivileged user to gain access to out-of-bound memory, leading to an integer overflow and a system crash or a leak of internal kernel information.



Version 1.1.0 bug fixes

The following issues have been resolved:

Search, Dashboard and Indexing

- An issue where some wrong query strings or syntaxes did not trigger the proper error message.
- An issue where canceling a search process for simple queries failed to stop the processes that ran for a considerable time.
- An issue that could cause a search timeout while executing heavy queries.
- An issue where some dashboard widgets did not align properly when users tried to arrange manually. Additionally, a horizontal scroll bar was seen in some dashboard widgets, even when the widgets were expanded to fit the content.
- An issue where the publicly shared dashboard widgets did not load while embedding in HTML pages in some browsers.
- An issue where the timechart query with a larger time range did not provide the complete result in dashboard widgets.
- An issue where a multi-line description of an alert was not completely displayed in the search results.
- In some cases, the audit logs for a search process stored an incorrect value in the **destination_address** field.
- While drilling down on a URL field of a search result, the field's value was truncated in the added filter of the search query.
- An issue where the *whois* option did not return proper results while drilling down an IP address in a search result.
- An issue which caused SLS to sometimes display partial results for key-value based search queries.
- In some cases, correlation queries could not be processed because of unhandled exceptions.
- An issue that sometimes caused the performance to degrade significantly in SLS machines with large repos.
- An issue where queries with a hyphen(-) in the **sls_name** field value did not return search results.
- SLS was unable to display correct results if users performed a drill-down operation on a field renamed by the *rename* command.
- An error occurred while saving a search if the search query contained regular expressions of the form `<string: all>`.
- An issue where many small-sized index partitions were created, which could block log collection.
- Selecting the *Sankey* chart in the search results of a query with a significant number of grouping parameters sometimes made the browser unresponsive.
- Sometimes, the donut chart rendered a blank panel until users changed the visualization option and switched back to the donut chart.
- In some cases, the visualization for *timechart* queries did not display the X-axis label up to the end of the selected timerange.



- An issue where SLS did not return search results when a search query contained the **repo_name** field and a NOT operator with a list.
- An issue where search became unresponsive in some cases, resulting in a search timeout.
- In some cases, SLS did not return search results while using aggregation and correlation queries requiring heavy computation.
- If a user changed the visualization for widgets in a shared dashboard, the change was not reflected for the other users. The visualization changed by a user with the *Edit* or *Full* permissions is now reflected across all the users. However, the visualization changed by a user with the *Read permission* is only reflected for the current user.
- In some cases, a vendor dashboard activated by a user was unable to display the data in the selected repos.
- An issue where SLS did not return search results when the chart command was used with a join query.
- An issue that could slow down the search process and cause search timeouts in some cases. SLS has now optimized the thread synchronization to improve the search and reduce search timeouts.
- An issue where SLS ignored the ampersand (&) symbol used in the search query and displayed matching results for only the remaining characters.
- An issue where SLS displayed empty results when users performed drill-down on the search results of queries with time functions such as minute, hour, or day as grouping parameters.
- The values of the **log ts** field were not displayed in the correct time format while using the *distinct_list* command.
- An error occurred when users searched for a query containing the % character from their search history or a widget.
- An issue where an empty dynamic list used in the search query with a non-existent field returned logs in the search result.
- An issue where using the timechart command with the *count()* and *distinct_list()* functions in a search query displayed values in empty time buckets. Additionally, the aggregation functions with the timechart command now display empty results instead of 0 in case of no events.

Alerts and Incidents

- In some cases, when the *Flush On Trigger* option was enabled for alert rules having correlation queries, SLS did not trigger alerts, or generated incidents from the previously triggered incident logs.
- When invalid **toList** or **toTable** process commands were used in the alert or dashboard, a gradual increase in resource usage occurred.
- The alert throttling feature only worked if the last set of values matched the current set of values. SLS now checks all the values within the throttling period to resolve the issue.
- In some cases, alerts failed to trigger when SLS failed to process multiple concurrent config regenerations.
- An error occurred if the Jinja variables were not placed within quotes in the body while configuring HTTP notifications of alerts.
- An issue where the integers in the query string field of HTTP notifications of alerts were converted to strings.
- The basic authentication header for HTTP notifications of alerts was unnecessarily encoded.



- In some cases, the **Content-Type** header was not set to JSON for HTTP notifications of alerts.
- The query string for HTTP notifications of alerts was converted to JSON only if the body was not empty.
- An issue where alerts were not triggered and users' dashboards with the *Pre compute* option disabled were not populated when the system was configured with alert rules having correlation queries requiring heavy computation.
- In some cases, aggregation queries in alert rules with a user-defined search interval could not be scheduled while reactivating the alert rules.
- An issue where Incident API endpoints requested without mandatory JSON parameters or proper JSON formatting were not handled correctly with a relevant error message.
- An incorrect notification message was displayed while reassigning incidents.
- Timestamp data in the *alert incident* data panel did not display according to the format defined in the Jinja *datetime* filter. While including timestamp data in the Jinja template, it is now mandatory to include the *datetime* filter.
- An issue where alert rules cloned from shared alert rules were not editable.
- An issue where alerts did not trigger when the chart and timechart commands were used together in the alert rules.
- An issue where an alert rule did not generate new incidents if they contained the same data compared to the incidents generated in previous time intervals.
- An issue where SLS did not send HTTP notification for alert rules when the alert rules had a newline character in the *Description* field, and the *Body* field in the HTTP notification configuration had a Jinja syntax for the *Description* field of the alert rules.
- An issue where alert rules did not trigger due to the non-decodable multipart format in the incoming logs.
- In some cases, alert rules with alert throttling enabled did not trigger even after the specified time.
- An error occurred when users tried to view data of the incidents generated with alert rules containing the % character in the query.
- An error occurred while saving the changes in cloned vendor alert rules having **col_ts** field in its Jinja template.
- An issue where incident data fields were not displayed correctly while viewing data of the incidents generated from the alert rule containing timechart query.
- An issue where notifications settings configured for an alert rule were not exactly replicated when the alert rule was cloned.
- An issue where SLS displayed random characters in the search results if the alert rules had a newline character in the Description field, and the Message field in the SNMP notification configuration of the alert rules had the `{{description}}` Jinja placeholder.
- An issue where incident ownership of alert rules was not changed when the *change incident ownership* action was performed in bulk for the alerts rules.
- An issue where SLS failed to update configurations of email notification for all the multiple alert rules selected in bulk.
- An issue where SLS failed to trigger alerts due to authentication failure in the database.

Reports

- An issue where a shared report templates became unshared when a different user edited it.
- An issue where integers were displayed as floating point numbers in reports.



- While deleting report templates, the **documents** field in the audit logs had an encoded value instead of the report name.
- In some cases, the colors displayed in the footer of report templates and the treemap representation of dashboard widgets were different than the colors chosen by the user.
- An issue where some Unicode characters were not displayed in the scheduled reports generated from report templates, and ad-hoc reports generated from dashboards, search, or search templates. Users can now select a different font while creating reports from layout templates to support Japanese Unicode characters in scheduled reports.
- An issue where users with the & character in their username could not view or download reports.

Licensing

- An issue where users could not edit the device IP when the number of devices reached the licensed limit.
- An issue in the SLS UI when a license with an invalid date was imported.
- An issue where users were unable to view the details of a SLS Collector's license.

System Monitor

- An issue in the System Monitor's line graph when users selected a date format other than *YYYY/MM/DD*.
- An issue where the CPU Usage chart displayed old results. The issue has now been resolved. However, a delay of at most 5 minutes can exist.
- The SNMP OID to display the statistics for service disruption returned an error when no services were disrupted.
- The *Disk Usage* chart in the *System Monitor Dashboard* sometimes displayed incorrect values.
- An issue where all the charts in the *System Monitor Dashboard* ended with the value of *zero*.
- An issue of high memory usage by a few services in the system.
- The SNMP OIDs to display the statistics of ZFS pools did not return the expected results for SLS instances having Ext4 filesystem and using ZFS-based storage.
- The SNMP OID to display the log queue in the collection process returned 0 even when there were multiple logs in the queue.

User Configuration

- An issue where user groups were not editable because they contained a reference to non-existing users in some cases.
- An issue where duplicate permission groups were not replaced while importing SLS configurations.



Repos and Storage

- An issue where repo names did not support a number at the start.
- An issue where high availability repo logs were buffered for a longer time, resulting in high disk utilization.
- In some cases, users could not create HA repos in multiple Distributed SLS instances from a SLS machine.
- When users attempted to change the repos of an activated vendor dashboard, *All Repos* was selected instead of the repos chosen by the users.

Authentication

- An issue where SLS displayed distinguishedName associated with the LDAP user irrespective of the Jinja placeholder defined in the *Username template* while configuring the LDAP strategy.
- The *Change Password* option is now only available for the users logged in using SLS Authentication.
- An issue where users authenticated through authentication plugins could not re-verify their username and password. As a result, the users could not perform activities that required password verification, such as activating or deactivating other users. From now on, the users authenticated with SLS's basic authentication, and the LDAP and Radius authentication plugins can re-verify their username and password.

Search Templates

- In some cases, the widgets in a search template were unable to populate the selected visualization.

Other Bug Fixes

- The message shown when a user is locked out now mentions *SLS Administrator*.
- Stormshield logo has been added in the front and back pages of generated reports. Additionally, the color of the charts and the footer now matches Stormshield's color scheme.
- The color scheme in various UI components and charts has been updated to match Stormshield's color scheme.
- An issue where the browser tab title did not accept the hyphen [-] character.
- An issue where the *Resolve All* option in the Notification Centre resolved only one notification.
- An issue where exporting logs failed when the SSH connection between the server and the remote host remained idle for too long or was disconnected.
- An issue where SLS packages did not sort as per their install date when users set the date format without a year at the beginning.
- An issue where users could not edit the filename pattern of an SCP Fetcher configuration if the value matched with the filename pattern of an existing configuration.
- An issue where SLS displayed an incorrect number of devices while importing configurations from *Settings >> System >> Sync*.



- Users belonging to the child groups in an Active Directory were unable to log in using LDAP Authentication if the corresponding *User Base DN* did not match the *Group Base DN*.
- An issue that sometimes caused high memory usage while opening the *Updates* page.
- An issue where SLS unnecessarily checked for operating system updates daily.
- Users were unable to update the NTP configurations of a SLS machine from the UI if the existing configuration file did not contain any server configurations.
- If multiple SSH notifications were triggered within a small interval, SLS was unable to send all the triggered notifications.
- An issue where unnecessary log files filled up the */opt* partition, resulting in high disk utilization.
- An issue where old files were not cleaned up properly, resulting in high disk utilization.
- While sending syslog messages over TLS to SLS, an active TLS connection did not shut down properly even when the client sent an alert to disconnect.
- An issue where a report template export file stored the email address configured in the template as well.
- An issue where high availability repo logs were buffered for a longer time, resulting in high disk utilization.
- An exception that could cause issues in log storage and search in some cases.
- An issue of log collection with the Syslog collector having multiple TLS connections from the same device.
- An issue where some critical information could be exposed in some process monitoring tools like *htop* and *pspy64* while creating and restoring configuration backup.
- Some unnecessary warning logs seen in the system have been handled.
- An issue where multiple unnecessary processes caused a heavy load to the system.
- An issue where expanding the *Create Alert* and *Create Dashboard* dialog boxes created additional blank space and downsizing the dialog boxes hid the content.
- An issue where the user-created ZFS pools were not imported automatically during the boot in the system with NVME SSD disks or multipath disks. The issue has been resolved by importing the pools from the *zpool.cache* file. However, the pools that are not present in the *zpool.cache* file need to be imported manually.
- An issue in the auto-tuning feature led to some performance issues in the system. The auto-tuning feature has been enhanced to proactively tune necessary resources allocated for SLS services according to the resource availability.
- While importing devices from a CSV file, an error occurred if users provided the IP address of a Syslog forwarder in the **distributed_collector** field.
- If users added a new repo or a device group, the permissions for a user group with the *All Repos* or *All Device Groups* permissions were not updated to include the newly added object.
- Alert HTTP notifications did not support the *incident_id* Jinja placeholder variables.
- A validation error occurred while configuring notifications for alert rules if users used *extra_info.query* in the Jinja template.
- An error occurred in search when the fields ending in *_ts* were used as parameters in all the aggregators except *distinct_list*.
- An issue of log collection with the Syslog collector due to the TLS connection shutdown after the idle timeout of 120 seconds.



SLS 1.0.0 key features

Stormshield Log Supervisor is a log management solution. It collects streaming data coming from Stormshield Network Security firewalls, analyzes it, and provides meaningful insights to your data in real-time. SLS allows tight control over widely distributed enterprise networks from a single location and offers capabilities of synthesizing the underlying risks associated with complex distributed attacks on large networks.

Stormshield Log Supervisor can be deployed on two virtual environments:

- VMware ESXi with the Stormshield Log Supervisor OVA,
- Microsoft Hyper-V with the Stormshield Log Supervisor Hyper-V VHD.

For more information on SLS v1.0.0, see the *User Manual* and our *Deployment Guides for OVA or Hyper-V VHD* on the [Stormshield Technical Documentation](#) website.

The version 1.0 of SLS provides the following key features.

Stormshield products

- The Stormshield Repo, Routing Policy, Normalization Policy, and Processing Policy have been preconfigured to ease the addition of new Stormshield Network Security devices.

Dashboards

- Stormshield dashboards have been enabled by default.
 - SNS Overview
 - SNS Threats
 - SNS Data
 - SNS Web
 - SNS Applications
 - SNS Hardware
 - SNS System
- SLS allows to modify default dashboards and to create new dashboards.

Search

SLS collects logs and stores them securely until the time specified in the system. You can search these logs using the SLS query language. You can further use the search results to create dashboard widgets, generate reports, and add alert rules. Using these features, you can monitor various compliance requirements, configure different correlations, and act on the incidents which require a prompt response.

- Search results can be filtered and drilled down to provide a defined range and log specifications.
- SLS paginates the search results for simple search queries. The pagination bar at the bottom lets you select a page number and the number of logs to show in each page. SLS only retrieves a maximum of 10000 logs for paginated results.



- SLS's **Query Language** is extensive, intuitive, and user-friendly. It covers all the search commands, functions, arguments, and clauses. You can search the log messages in various formats depending on the query you use. The query has been categorized as **simple search, aggregators, one-to-one commands, process commands, filtering commands, and pattern finding commands**.
- SLS supports the chaining of commands and multi-line queries. Use a pipe (|) to chain the commands. The query bar automatically adds a new line after you reach the end of a row. Alternatively, you can also press Shift+Enter to all lines.
- A **Search Template** stores search queries with placeholders. The stored queries are called **base queries**. The placeholders in the base queries are used as variables which you can replace with actual search keywords during runtime. You can drill down any field value in the search results directly into a search template. This helps you investigate the corresponding field-value pairs in the search templates easily and efficiently.
- The following features are available in **Search Templates**:
 - You can add multiple base queries to a Search Template.
 - You can add multiple dashboards to a Search Template View.
 - The default time range of a Search Template Widget can be overridden from the Search Template View.
- You can use a **macro** to save any search query in the system and use it in the Search, Dashboards, Reports, Alert Rules, Label Packages, Search Packages and Search Templates. You can compress long queries into a single name and re-use it in the system. You can use a macro in different settings and update from the macros page in Settings >> Knowledge Base. You can also import or export macros.
- **Aggregation functions** are used with the **chart** and the **timechart** commands to aggregate the fields. You can use the aggregation function *distinct_list* to view the list of the distinct values of a particular field.
- **Search Views** enables you to first define table-views and then use it to view the result in a structured manner. You can also view the Top-10 statistics for the selected fields in the Search Views Interface and access the previously created Search Views from the Search Landing Interface.

Dynamic List and Tables

- A *dynamic list* collects the specified field values during the runtime and stores them for a limited or an unlimited period. You can populate or update the dynamic list using the process command *toList*. You can also update the list items via alerts and dashboards. This feature reduces your manual entry, providing you with a real-time updated list.
- A *dynamic table* stores the specified field and field values during the runtime for a limited or an unlimited period. You can populate or update the dynamic table using the process command *toTable*. You can use the dynamic table as an enrichment source to enrich your logs.



Interesting Fields

- The Interesting Fields are the relevant fields presented based on the data distribution for the results of Simple Search queries, One-to-One commands, and Filtering commands. The following parameters are used to measure the data distribution:
 - Percentage [%] displays the percentage of occurrences of a given field.
 - Count [# of values] displays the number of unique values for the given field.
 - Mean Deviation displays the value of the deviation from the average of occurrences of the given fields.
 - Median Deviation displays the value of the deviation from the fields with the highest number of similar occurrences.

Data Privacy

- SLS incorporates the **Data Privacy Module** feature, which functions under the Four Eyes Principle. The Data Privacy Module demands a certain activity to be approved by at least two people, thus facilitating the delegation of authority and increasing transparency within the organization. In SLS, the users with the SLS Administrator privilege can selectively encrypt the field values. These field values are encrypted in the search, dashboard, alerts, incidents, and reports.
- To decrypt the field values, a user must create a data privacy request, and another user with granting permissions must approve the request. This way, any confidential data will remain encrypted unless agreed otherwise by two users.

Reports

A SLS report is the collection of information, events, and findings which are collected, analyzed, and presented in an organized manner. You can view all the generated reports, rules to generate the reports, and the report templates in the **Reports** page.

- The generated reports are populated in the inbox along with their names and their corresponding formats (PDF, XML, HTML, DOCX, or CSV).
- SLS generates audit logs indicating the success or failure of reports generated for each query. In case of failure, the logs also contain the reason of failure. Additionally, the success or failure status is also displayed in the UI.
- You can easily free up memory and remove clutter in your report inbox by deleting all reports older than a selected date or by deleting all the reports within a specific date range.

Alerts and Alert Notification

Alerts in SLS are warnings generated to notify users when any significant events occur. They fire **incidents** that enable you to execute appropriate actions. Any valid search query can trigger an alert to generate incidents.



- The alert rules support the following features:
 - You can prevent the same set of values from generating the same incident continuously. While creating an alert rule, enable the **Alert Throttling** check box and select a field and a time interval. The alert rule does not fire for the same set of values of the selected field until the specified time interval.
 - You can specify the details of the events in the Jinja format while creating an alert rule and view the corresponding details in the Incident Data panel of the generated incidents.
 - While creating an alert rule, SLS allows you to define the frequency in which the search should be performed using the provided query.
- You can create an alert rule and select the incident notification medium. SLS can notify you via e-mail, SSH, SNMP, HTTP, or Syslog. You have the flexibility to design an alert mechanism based on your requirements.
- You can use the Jinja template `{{extra_info}}` while adding the alert notifications for *Email Notification*, *SNMP Notification*, and *Syslog Notification*. This allows you to view the values of the field `extra_info`, which contains information about the queries and repos.

Settings and Configuration

- SLS supports multiple data paths for the storage devices connected to the system.
- The Routing Policy allows you to selectively determine what incoming data gets forwarded to a particular repository and what gets dropped. Using this feature, users can choose to:
 - store raw message along with normalized data
 - discard the raw message and store only the normalized data, or
 - discard the entire event (both the raw message and its normalized data).
- The auto-tuning feature automatically tunes the SLS backend components such as premerger, merger, index searcher, analyzer, file keeper, and normalizers based on the load on these components. A li-admin command *auto-tuning* is exposed to start and stop the tuning process. The auto-tuning feature is **enabled** by default.
- You can configure SLS to allow access to only those server names and aliases specified in the HTTPS certificate, ensuring that SLS does not respond to arbitrary host headers.

Login and Lockout Policy

- You can configure a unique ID field for LDAP users from the LDAP strategy. The field can uniquely identify each LDAP user, eliminating the need to reconfigure the user and their personalized items if you move them to another group within the Active Directory.
- SLS implements the account lockout policy, which locks a user account for a specific duration if users make multiple failed login attempts. A SLS administrator can set the values for the lockout threshold and the lockout duration.

User Account Management

SLS provides you an intuitive interface for user account management and authorization under the *User Accounts* section. It lets you configure different permissions into a permission group. You can then assign the *permission group* to respective users in a *user group* based on their roles and responsibilities. This sequential process enables role-based access control over your system. You can also assign the *user groups* to *Incident User Groups* and *Data Privacy Groups*.



- You can better support the workflows of security analysts and the general processes within organizations using Incident User Groups. These groups map to SLS User Groups and allows to reassign incidents.
- You can select the permissions to *read*, *create*, and *delete* label packages while creating permission groups.

Backup and Restore

Backup and Restore settings enable you to create backups of repos and configurations and restore them in the future. You can use any SFTP client to download and upload the backup files.

- You can create two types of backups in SLS, named *Configuration backup* and *Logs backup*.
 - *Configuration backup* stores the system-related information such as users, devices, permissions, policies, and repos.
 - *Log backup* stores the logs and indexes.

System Monitoring

- You can easily and effectively monitor your SLS system using system monitoring. You can retrieve information about multiple components, services, and processes without logging into the SLS UI.
- SLS has implemented multiple mechanisms to avoid data corruption and performance degradation when the disk storage is full. SLS notifies users when the disk usage reaches a specific threshold limit and stops collecting logs if there is not enough disk space.



Contact

To contact our Stormshield Technical Assistance Center (TAC):

- <https://mystormshield.eu/>
All requests to technical support must be submitted through the incident manager in the private-access area https://mystormshield.eu, under Technical support > Manage cases.
- +33 (0) 9 69 329 129
In order for us to provide high-quality service, you are advised to use this communication method only to follow up on incidents that have been created earlier on https://mystormshield.eu.



STORMSHIELD

All images in this document are for representational purposes only, actual products may differ.

Copyright © Stormshield 2023. All rights reserved. All other company and product names contained in this document are trademarks or registered trademarks of their respective companies.