GUIDE

STORMSHIELD ENDPOINT SECURITY EVOLUTION

ADMINISTRATION GUIDE

Version 2.7.2

Table of contents

In the documentation, Stormshield Endpoint Security Evolution is referred to in its short form: SES Evolution.

1. Getting started

Welcome to the Stormshield Endpoint Security Evolution administration guide version 2.7.2.

This guide contains all of the essential technical information to run and monitor the product in your environment.

SES Evolution is a global security solution that offers comprehensive workstation protection in organizations of all and server sizes. The SES Evolution agent is installed on the workstations and the servers and transparently protects them from known and unknown attacks and intrusions. The agent is configured in an administration console and is in constant contact with SES Evolution agent handlers that distribute security policies.

In the administration console, users can also configure security policies and view event logs generated by workstations and servers, making it possible to monitor their operation.

SES Evolution also has a public REST API. It is not enabled by default. For more information on the public API, refer to the section Enabling and managing SES Evolution's public API.

2. Connecting to the SES Evolution administration console

1. Connect to your workstation using your Microsoft Windows domain account.

2. Run the Stormshield Endpoint Security Evolution administration console.
    You are now connected to the console with your Windows account.
    If the Windows account is not recognized or if the backend component cannot be reached, a connection window appears but the administration console does not open.

3. To see your account preferences, click on the icon in the upper banner in the console, next to your user name. The administration console appears by default in the language of your operating system; you can however change the language on the console.

To connect with an account other than the one for which the Windows session was opened, the executable file of the console can also be launched using the option **Run as another user**.

3. Understanding the dashboard

The SES Evolution dashboard provides an overview of the security status of your pool and how it is managed. You can identify the elements that cause issues at a glance, and access the various control or monitoring panels via shortcuts. It consists of two tabs:

- The **Security** shows indicators on events that occur on agents.
- The **Operation** tab shows information on operating conditions.
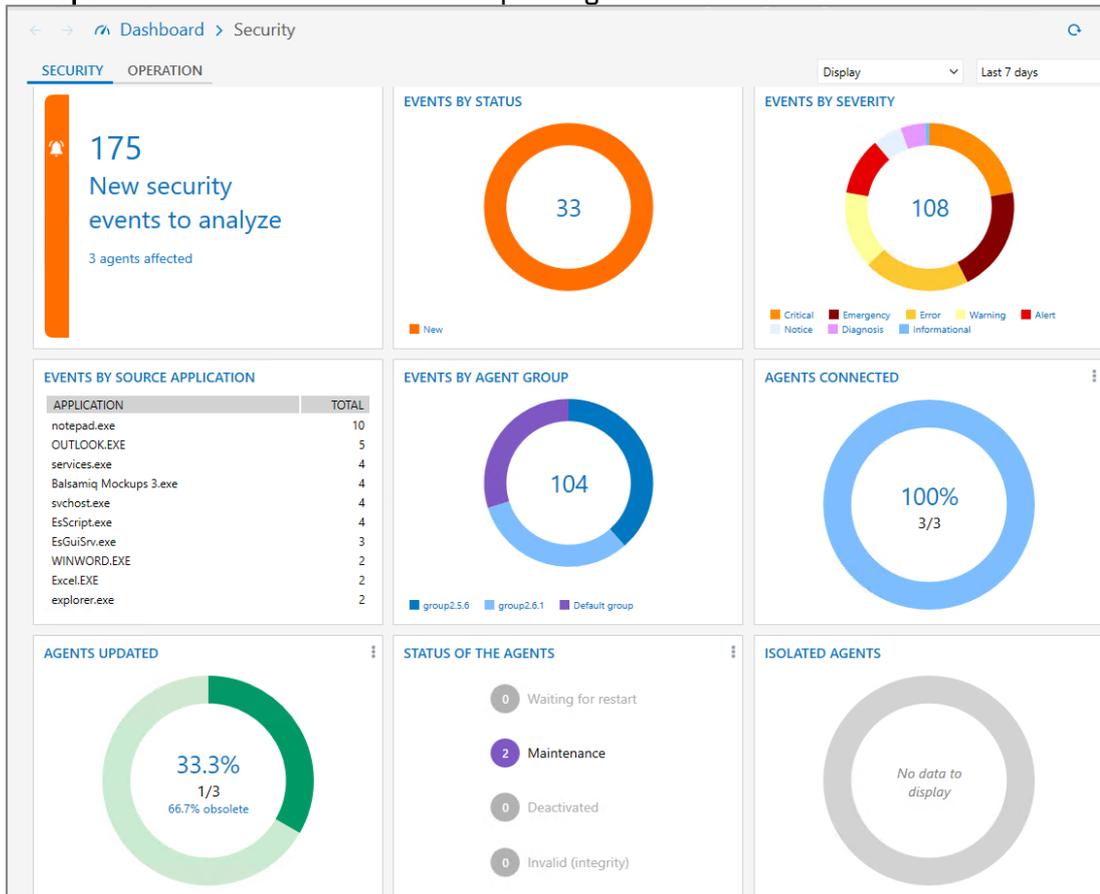
3.1 Monitoring SES Evolution agents

The **Security** tab in the dashboard provides a quick overview of the most critical events that occur on your pool, and the status of SES Evolution agents.

3.1.1 Monitoring events that occur on agents

An event is a grouping of several identical logs generated on several agents.

- The **New security events to analyze** tile shows the following information without time limit:
  - Number of *New Emergency* and *Alert* events,
  - Number of agents affected,
  - Number of events that appeared over the past 24 hours,
  - Percentage of the latest events that appeared over the past 24 hours, out of the total number.

  Click on the number to see the list of *New Emergency* and *Alert* agent logs.

- By default, the other **Events...** tiles show indicators over the past 24 hours. You can modify this period from the drop-down list in the top right corner.
  Click on the title of the **Events...** tiles or on the various parts of the diagrams or lists to go directly to the **Environment > Agent logs** panel and see the list filtered by the selected period and type of indicator.

3.1.2 Checking agent status

- The **Agents connected** tile shows the number of agents that reconnected to their agent handler within the normal time defined for their group.

- The **Agents updated** tile shows the number of agents whose software, policy and configuration versions match those defined by their agent group. They may sometimes have a higher software version if it cannot revert to an older version and if these agents were forced to update.

- The **Agent status** tile shows whether agents are in one of the following statuses:

Status	Description
Waiting for restart	The agent had to be restarted to complete an installation, an update or to apply changes.
Maintenance	Maintenance mode is enabled on the agent.
Disabled	The agent was disabled by a challenge.
Invalid	The agent reported issues after an integrity check.

Click on the circles, text and numbers of agents to go to the **Agents** panel and view the filtered list.
Click the tile titles to go directly to the general panel for agents. For more information, please refer to the section Viewing agents in the console.

Click ⋮ to export the list of all agents in the pool, or a list of agents by status, to a .CSV file.

- The **Isolated agents** tile shows the number of isolated agents on the network and their isolation status:

Isolation status	Description
Isolated	Agent that was isolated from the network
To be isolated	An isolation was requested for the agent, but the agent has not run it yet.
Undoing isolation pending	A request to undo isolation was submitted for the agent, but the agent has not run it yet.

Click on each part of the chart to go directly to the **Isolate** panel and view the filtered list.
Click the tile title to go directly to the general isolation panel. For more information, refer to the section Isolating computers from the network.

3.2 Monitoring operational resources

The **Operation** tab in the dashboard provides a quick overview of the status of backoffice resources and operating conditions.

3.2.1 Finding out an SES Evolution agent's version

The **Agent versions** tile displays the distribution of software versions in your pool. Scroll over the part of the circle corresponding to a version to show the number of agents concerned.

Click on the circle, text and numbers of agents to go to the **Agents** panel and view the filtered list.

Click the tile titles to go directly to the general panel for agents. For more information, please refer to the section Viewing agents in the console.

3.2.2 Monitoring licenses

The **Licenses** tile shows license information in a diagram.

The diagram will show the number of active agents and the proportion compared to the number of agents allowed in the license. An agent is considered active if it has connected to the agent handler within the past 10 days. The color of the diagram changes according to the proportion of licenses used.

Green	The number of active agents is below 90% of licenses' full capacity.
Orange	The number of active agents is between 90% and 110% of licenses' full capacity.
Red	The tolerated threshold of 110% has been exceeded.
Gray	The license has expired.

License information is refreshed every hour and every time you access the dashboard.

Click on **Licenses** at the top left side of the tile to go to the panel for licenses. For further information, refer to the section Managing SES Evolution licenses.

3.2.3 Checking server status

The **System** tile shows the statuses of various servers in different colors: backend server, databases and agent handlers. For more information, refer to the *SES Evolution Installation Guide*.
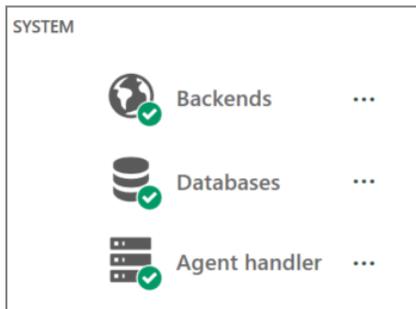
If an issue is detected on a server, SES Evolution will generate a system log, and the color of the server's icon changes.

Click on the icons and names of servers to go to the **System logs** panel and view the filtered list.

Backends

The backend is the application server that centralizes all operations performed in the SES Evolution environment. The backend icon changes color according to the amount of resources consumed:

Green	All backends are running,
Orange	The average RAM or CPU consumption of one or several backends exceeds 90% (moving average over one hour), or the disk space used exceeds 75%.
Red	The status of one or several backends has not been updated for more than 5 minutes or the task of deleting logs was not successfully carried out, or the disk space used exceeds 85%.

Click on ⋯ to obtain more accurate information on each backend's resource consumption and the date of its last connection. This information can also be accessed from the upper banner of the console.

The result of the log deletion task is also shown. If the task failed, move your mouse over the red cross of the **Task** column to show the exact error message. For more information on this task, see the section Monitoring databases.

Databases

SES Evolution runs with several databases, including an administration database and one log database.

The color of database icons 🗄️ changes according to whether they can be reached, and the amount of disk space used:

Green	All databases can be reached and less than 70% of the disk space is used. Saturation of the database is estimated in more than three months.
Orange	On at least one database, between 70% and 80% of the disk space is used, or saturation of at least one database is estimated in one to three months.
Red	At least one database cannot be reached, or the hard disk space used exceeds 80%, or the saturation of at least one database is estimated in less than one month.

Click on ⋯ to see the percentage of disk space used and when each database last connected. This information can also be accessed from the upper banner of the console and via the **Backoffice > System** menu. For more information, see the section Monitoring databases.

Agent handlers

The agent handler receives data and logs directly from agents, and updates the administration database via the backend. The color of the agent handler icon 🖥️ changes according to its resource consumption:

Green	All agent handlers are operational.
Orange	The average RAM or CPU consumption of one or several agent handlers exceeds 90% (moving average over one hour), or the disk space used exceeds 75%.
Red	The status of one or several agent handlers has not been updated for more than 5 minutes, or the disk space used exceeds 85%.

Click on ⋯ to obtain more accurate information on each agent handler's resource consumption and the date of its last connection. This information can also be accessed from the upper banner of the console.

3.3 Sending dashboard indicators by email

You can configure SES Evolution to send activity reports by email to recipients of your choice. The activity report will contain all dashboard indicators.

You must first configure an SMTP server. For more information, refer to Configuring an SMTP server.

You must hold the **Email Notifications-Modify** permission to configure the sending of activity reports.

To send dashboard indicators by email:

1. In the **Backoffice > System** menu in the administration console, go to the **Email Notifications** tab.

2. Click on **Edit** in the upper banner.

3. In **Activity reports**, click on **Add rule**.
   The rule creation wizard opens.

4. Enter the **Rule name** and **Frequency** with which you wish to send activity reports, and click on **Next**.

5. In the field at the bottom of the screen, enter the email address of the user who will receive activity reports, select the language, then click on **Add**.

6. Add more email addresses if you wish to send reports to several recipients.

7. Click on **Create**.
   The rule will be added to the table in **Activity reports**.

8. Add other rules if necessary.

Reports are sent by email at 00:00 every day, every Sunday, or every last day of the month, depending on the chosen frequency. The subject of the email generally looks like this: *SES Evolution - Activity report from 11/12/2023 to 11/12/2023*.

You can disable or enable rules again by clicking on the checkbox in the **Enabled** column. The action buttons to the right of a rule can be used to duplicate or delete it.

You can temporarily stop sending activity reports by disabling the **Enable notifications** option.

With SES Evolution, you can also send alerts by email based on the types of logs generated. For more information, see Sending agent logs alerts by email and Sending system log alerts by email.

4. Managing SES Evolution licenses

You have registered a license while installing your SES Evolution environment.

Licenses determine the number of active SES Evolution agents that you can manage with the solution, and have an expiry date.

Several licenses can be imported, in which case, the number of agents allowed is the total number of agents for all licenses.

4.1 Importing the license in SES Evolution

You must hold the **Licenses-Modify** privilege to be able to import licenses.

1. In the **Operation** tab in the administration console dashboard, click on **Licenses**.
2. Click on **Add a license** and choose the license file (e.g., *SES-JCCA-WE9T-Q5RA.lic*). The **Capacity** field represents the number of active SES Evolution agents and the total number of agents allowed per license,

4.2 Reading license information

You must hold the **Licenses-Display** privilege to be able to read license information.

The **Operation** tab in the **Licenses** section of the administration console dashboard shows the number of active agents and the proportion compared to the number of agents allowed. An agent is considered active if it has connected to the agent handler within the past 10 days. The chart is green when the number of active agents is below 90% of the license's full capacity, orange when it is between 90% and 110%, and red when the tolerated threshold of 110% has been exceeded.

License information is refreshed every hour and every time you access the dashboard.

5. Managing users on the SES Evolution administration console

Users access the console with their Microsoft Windows accounts which must be on the same Active Directory domain as the backend component. If this is not the case, then a relationship of trust must be established between the domains.

By default, only the super administrator specified during installation can log in to the administration console. He/she can then add other users or groups of users who will be able to log on in turn.

> **ⓘ NOTE**
> If you rename the Windows account of this super administrator, make sure that you have created a user beforehand with the new name in the SES Evolution administration console. Otherwise you will not be able to log in to the console. For further information, refer to the section Adding users on the administration console.

Each user or group is assigned a role that defines their profile and restricts the functionalities available in the administration console. Three roles are available by default: Audit, Helpdesk and Administration. New roles can also be created and customized.

> **✎ EXAMPLE**
> You can create an *AdministratorsSES Evolution* group in Active Directory, add it to SES Evolution, and assign it the *Administration* role. In this case, all users in the group will automatically have administrative privileges in the SES Evolution console. You will not need to add them individually.

Multiple users can log on simultaneously to consoles managing the same pool.

5.1 Creating custom roles

Only a user with the *Administration* role is authorized to add other users.

1. Choose the **Backoffice > Users** menu, then the **Roles** tab.
2. Click **Edit** in the top banner, then click on **Create role**.
3. Enter a name for the role and its description if necessary.
4. Click on **OK**. The new role appears in the list. The most restrictive privileges are applied by default.
5. For each privilege, choose the type of access that you want to grant. Every privilege corresponds to a panel in the administration console. By default, only the panels **Deployment**, **Dashboard** and **Licenses** are accessible.
The **Lock** privilege makes it possible to break locks set up by other users on panels in the console. For more information on locks, see Managing simultaneous user connections to consoles administering the same pool.

5.2 Adding a user or group of users to the administration console

Only a user with the *Administration* role is authorized to add other users or groups.

1. Choose the **Backoffice > Users** menu, then the **Users and groups** tab.
2. Click on **Edit** in the top banner, then on **Add** in the **Users** or **Groups** area.
   An empty line is displayed.
3. You can:

   - Click on the ☐ icon to the right of the line to select a user/group in the Active Directory.
   - Manually enter an Active Directory user name/group using the syntax *DomainName\samAccountName* for a user, and *samAccountName* for a group.
   - Enter a local user name manually.

   SES Evolution checks the validity of the user/group and displays its status on the right.
   Hover your mouse over the icon in the **Status** column to obtain more information.

4. Select the role to assign to the user/group:
   - **Audit**: this role makes it possible to view all panels in the console and edit the settings of the user's own account, but no other modification and deployment operations are possible. This role is dedicated to log reading and agent monitoring.
   - **Help desk:** This role holds the same privileges as the Audit role. In addition, it allows the user to respond to challenges and unlocks locked operations. This role is dedicated to the maintenance of the SES Evolution pool.
   - **Administration**: This role makes it possible to perform all operations accessible in the administration console without restrictions.
   - **Custom role**
5. In the **Group** area, sort the groups by priority using the arrows in the **Order** column. If a user belongs to more than one group, he or she is assigned the role of the group with the highest priority.

   If a user is declared individually AND via a group, the individual user role is assigned.

5.3 Managing the simultaneous connection of users to consoles that manage the same pool

Multiple users can simultaneously manage the same pool from different hosts.

When a user edits any of the following resources, they will automatically be locked and no other users can edit them:

- An agent group,
- A policy,
- A rule set,
- A Yara or IoC analysis unit,
- An agent handler group.

> ✏ **EXAMPLE**
>
> - While user 1 is modifying policy A, user 2 cannot modify it, but they can modify the rule sets contained in policy A, and policy B.
> - While user 1 edits agent group A, user 2 cannot edit it, but may add a new agent group B.

When a user saves or cancels changes, the panel will automatically be unlocked if there are no more objects being edited in this panel.

No user can edit a panel locked by another user. The **Edit** button in the upper panel is replaced with a padlock. Hover the mouse pointer over the padlock to see who locked the panel and since when. There are three types of padlocks:

Padlock	Description
🔒 Green	You are editing the resource and it is locked for the other console users. Only users with **Locks - Unlock** privileges can unlock the resource. In this case, you are informed when you save your modifications.
🔒 Orange	The resource is being edited by another user, but you have **Locks - Unlock** privileges and you can thus unlock the resource to edit it. First ensure that your action is legitimate.
🔒 Red	The resource is being edited by another user, and you do not have **Locks - Unlock** privileges. Therefore, you cannot unlock the resource.

Unlocking a resource can be particularly useful when it accidentally remains in edit mode for example.

As this operation releases the panel and cancels the other user's changes in progress, it must be used carefully. In this case, the user who held the lock first will be warned when s/he attempts to save changes.

To unlock a panel if you have the appropriate privileges, i.e. if the padlock is orange:

1. Click on the padlock in the upper banner.
2. Confirm the operation in the window that appears.

6. Configuring SES Evolution agent handlers

Agent handlers are SES Evolution servers that make it possible to distribute security policies and software updates to agents. It is also possible to receive:

- agents' event logs, which can be saved and sent to Syslog servers.
- the status of agent monitoring data and display them in the **Agents** tab of the agent group.

Every agent handler belongs to a group of agent handlers.

The parameters of each agent handler and each group must be defined. **Agent handlers-Modify** privileges are required.

We recommend that you install at least two agent handlers per handler group, installed on two separate servers, to ensure service continuity. For more information, see Ensuring Service Continuity in the *Installation guide*.

The following limits apply in order to prevent disk saturation of an agent handler:

- a limit of 500 MB on folders named "InvalidPackages" located in the "Normal" and "Urgent" folders at the location "%programdata%\Stormshield\SES Evolution\Server\AgentLogs". These folders store log packages sent by the agents that the agent handlers cannot manage properly.
- a limit of 100 MB on the folder named "InvalidCertificates" located at the location "%programdata%\Stormshield\SES Evolution\Server". This folder stores the certificates of the agents considered invalid (revoked or expired).

When these limits are reached, the oldest files are deleted to free up half of the folder storage capacity.

6.1 Creating groups of agent handlers

A group of agent handlers consists of one or several agent handlers. When an agent must connect to an agent handler, the agent gives priority to the last handler that accepted its request. If the connection fails, the agent will randomly choose another handler from the group until its request is accepted.

Once an agent handler is installed, it will automatically appear in the **backoffice > Agent handlers** menu in the administration console. By default, it belongs to a group named *New Group (agent handler name).* By default, you can edit this group, create new groups or move an agent handler to another group.

Agent logs can be sent to different Syslog servers configured for each agent handler group. For example, configure several Syslog servers to receive logs of varying levels of severity or with different content formats.

The Stormshield Log Supervisor (SLS) log management solution can be used with SES Evolution. For more information, see section Configuring communication with a Stormshield Log Supervisor (SLS) server and the SLS documentation available on the Stormshield Technical Documentation website.

> **ℹ NOTE**
> If a Syslog or SLS server is unreachable, the agent handlers temporarily store the logs for a maximum of 24 h, provided there is sufficient disk space. Logs older than 24 h are deleted and replaced by new logs.

6.1.1 Creating new groups of agent handlers

1. Choose the **Backoffice > Agent handler** menu.
2. In the left panel, click on the **+** icon. The *New group* line appears.
3. In the **Agent handler group settings**, enter the **Name** of your agent handler group.
4. If you have SES Evolution agents in versions 2.6 and lower, enable the **Enable old communication protocol** setting. This setting enables these agents to communicate with the backoffice using TCP on port 17000, while agents in versions 2.7.1 and higher use HTTPS on port 443.
   Once your entire pool has been migrated to 2.7.1 or higher, it is recommended that you disable this setting to reduce the potential attack surface.

5. If you wish to send logs from agents in this agent manager group to Syslog servers, click on **Add server** and set the following parameters:

- **Address**: Enter the IP address or DNS name of the Syslog server.
- **Protocol**: Select the protocol for communication with the Syslog server. If you wish to encrypt the data exchanged, select TCP/TLS. In this case, the root certification authority and intermediate authorities of the Syslog server must be imported into the certificate store of each agent handler machine.
- **Port**: Enter the port number used for Syslog (*TCP 1468* by default). The TCP or UDP port numbers indicated here are allowed on the firewall of the workstation that hosts the agent handler, as well as on all network devices located between the agent handler and Syslog server.
- **Transfer type**: choose the parameter defined during the installation of the Syslog server.
- **Structured data**: Use this field to specify additional data to insert in the header of Syslog messages. To know the expected data format, refer to RFC 5424. You can add more than one data in the field. For example: [ABC param1="value1"][KEY@12345 param2="value2"].
- **Message format**: choose the message format:
  - simple text mode (like the messages displayed in the **Agent logs** menu),
  - raw JSON format containing all the technical data,
  - CEF format,
  - IDMEF format.
- **Message language**: Select the language if necessary.
- You can indicate a maximum message size in bytes.
- Choose the lowest log severity to send to this server.
- **Context details**: Choose the context level you wish to send to the Syslog server:
  - **None**: No context details are sent to the Syslog server, which only receives strong signals of an attack (i.e., alerts).
  - **Simple context detail**: Strong signals are sent to the Syslog server, along with creation and termination logs for processes running on the agent at the time of the attack and shortly after the first attack log.
  - **Full context detail**: All attack-related logs are sent to the Syslog server, regardless of the severity level selected above.
    Whether the agent handler receives the full context detail depends on the agent group configuration. By default, transmission is deferred, and the simple context detail is sent well before the full context detail.

  For more information on context details, see Understanding what makes up a context.

If you have configured at least one Syslog server, a Syslog server operation indicator will appear in the top banner of the console, after the environment has been deployed. It indicates the presence of any warnings or alerts, if any. Click on the indicator to display details for each server.

6. If you wish to move an agent handler from another group to your new group, select the handler and drag and drop it to the new group.

7. Click on **Save** in the upper banner.

> ℹ️ **NOTE**
> We advise against the use of UDP for communications with the Syslog server. The TCP protocol

can be used primarily for configuration and testing purposes, but in production we recommend using the TCP/TLS protocol.

6.1.2 Configuring communication with a Stormshield Log Supervisor (SLS) server

If you are using SLS, the Stormshield log management solution, complete the settings as follows in the **Syslog servers** section of the agent handler group settings panel.

Sending logs over TCP

1. Indicate the IP address of the SLS server.
2. Select TCP as the protocol.
3. Indicate port 601.
4. Select **Raw JSON** as the **Message format**.
5. Select **Non-Transparent-Framing** as the **Transfer type**.

Sending logs over TCP/TLS

The SLS server must have an X.509 certificate in PEM format.

1. Enter the *.crt* certificate and *.key* private key in the SLS administration console.
2. Import the SLS server's root certificate into the certificate store of each machine hosting an agent handler.
3. In the **Syslog servers** section of the SES Evolution administration console, configure the settings as follows:
   - Indicate the host name or IP address of the SLS server. It must match the address entered in the certificate. If you have used the server host name in the certificate, you can specify the corresponding IP address in the HOST file of the machines hosting the agent handlers.
   - Select TCP/TLS,
   - Indicate port 6514,
   - Select **Raw JSON** as the **Message format**,
   - Select **Non-Transparent-Framing** as the **Transfer type**.

In both cases, if you encounter issues receiving logs in SLS, refer to System logs in the SES Evolution administration console. Filter by the relevant agent handler. For more information, see the section Monitoring the activity of SES Evolution backoffice components.

> **ⓘ NOTE**
> The TCP port number indicated here must be allowed on the firewall of the workstations that host the agent handler, as well as on all network devices located between the agent handlers and SLS server.

6.1.3 Troubleshooting

A Syslog server is not receiving logs from agents:

- *Situation*: One of the Syslog servers configured in an agent handler group is not receiving any logs from agents.
- *Cause*: There may be an error in the TCP/TLS configuration of the Syslog server.

- *Solution*: First check that the Syslog server is running in TCP mode.
  If so, check the logs issued by the agent handler. If there is a TCP/TLS configuration problem, the agent handler issues a log that identifies the defective Syslog server and describes the possible causes.
  If you do not see this log, try restarting the agent handler to force the issuing of logs. Wait at least one minute after the restart.
  Depending on the log indications, then review the TCP/TLS configuration of the Syslog server.
  You can also check the minimum severity level of the logs you have set, so that they are sent to the Syslog server.

6.2 Configuring the parameters of agent handlers

After an agent handler is installed, it will automatically appear in the **Agent handlers** panel in the administration console. By default, it belongs to an agent handler group named *New Group (agent handler name)*.

1. Choose the **Backoffice > Agent handler** menu.
2. Select the agent handler from the left panel.
3. Click on **Edit** in the upper banner.
4. Change the default **Name** of this agent handler if necessary.
5. Click on **Save** in the upper banner.

7. Managing SES Evolution agents

The SES Evolution agent is installed on Windows workstations and servers to detect or protect against malicious attacks. As for the SES Evolution agent handler, it provides the security policy and applies the corresponding protections. Agents send their status and event logs to the agent handler as soon as these events occur, allowing you to track the status of your pool from the administration console.

The agent connects periodically to the agent handlers in the handler group assigned to it, with priority given to the last handler that accepted its request. If the connection fails, the agent will randomly choose another handler from the group until its request is accepted.

When the agent is not connected to a network or if none of its default or backup agent handlers can be accessed, it will run autonomously by applying the last known security policies.

The agent saves logs locally for the entire time it is disconnected from the network. When it reconnects, it sends its logs to the agent handler. Logs can also be exported to a *.cab* file and imported into the administration console to be read. For further information, refer to the section **Reading logs of offline agents**.

To find out which Windows versions are supported by SES Evolution, see the **Product Life cycle** document.

7.1 Creating agent groups

An agent group is an SES Evolution agent template that you deploy on all workstations that need to share the same configuration, especially within the same security policy. Any subsequent changes to the agent group configuration are applied to all agents in the group.

> ✎ **EXAMPLE**
> Separate agent groups can be created for the following cases:
> - Servers and workstations of users requiring different security levels,
> - Departments in the company that require customized security rules,
> - Laptops of mobile employees and desktop computers, etc.

After an SES Evolution agent is installed on a workstation,, it appears in the **Agents** panel of the administration console. It will be automatically placed in the agent group to which it belongs.

The **Agent groups - Modify** privilege is required to create and configure agents.

To create an agent group:

A default agent group is automatically created in the console, but you can create custom agent groups.

1. Select the **Environment > Agents** menu.
2. In the left panel, click on **Create a group**. The line *New group* appears.
3. In the **Agents** tab in the right panel, enter a **Name** for the group.
4. Configure the agent group according to your preferences in the **Policies**, **Scheduled tasks**, **Settings** and **Status and logs** tabs. You must select at least one policy.
5. Click on **Save** in the upper banner to save changes.

To create a new group, you can also duplicate an existing group. While a duplicated group keeps all the settings of the original group, it does not contain any agents.

1. Select the group to duplicate.
2. Click on **Duplicate** in the ≡ menu.

7.2 Configuring agent groups

See following sections to choose the security policies applicable to your agents, configure tasks and analyses, agent updates, log management, etc.

7.2.1 Applying security policies to agents

You must apply at least one security policy to every agent group. Several secondary policies can also be added, and will apply when certain conditions are met.

> ✎ **EXAMPLE**
> You can add a conditional policy for mobile users, which applies when some workstations are no longer located within the internal corporate network. You could also define a quarantine policy that applies as soon as an agent's health indicators reach unsatisfactory levels.

To apply one or several security policies to an agent group:

1. Go to the **Policies** tab in an agent group.
2. In the **Policies** drop-down list, select the main security policy to be applied to all the agents in the group.

> 💡 **TIP**
> A blank policy is offered in the drop-down list. When one is used, the protection of an agent group (except self-protection) can be temporarily disabled, for example for tests and troubleshooting.

3. If necessary, click on **Add a conditional policy**. For more information, refer to Adding a conditional policy.
4. By default, the **Switch policies to detection mode** option is enabled: the policy rules do not apply a block but generate a log. Disable this option to switch the policies to Protection mode.

Add a conditional policy

1. In the window of the new conditional policy, select the policy which will be applied under certain conditions in the **Policy** drop-down list.
2. Click on **Add a condition** and give the condition a name.

3. Click on **Add a test** and choose from one of the following tests:
   IP address
   Enter an IP address, address range or subnet and choose whether it needs to be within range or out of range for the test to be validated.
   You can define several ranges separated by commas, For example *172.16.16.0/24,10.10.0.0/16*.
   Reachable agent handler
   Enable this option to indicate that the agent must be able to reach the agent handler for the test to be validated.
   Ping
   Indicate the IP address or network name of the host that you want to reach using pings, whether the agent must be able to reach it for the test to be validated, number of tries, and frequency of tries.
   Result of custom script
   Click on ••• to add a script, and specify its path, arguments and where to run it. Indicate what its **Result** must be for the test to be validated. This result must correspond to an output code of the script.
   It is best to use **Local service** as this is an account with restricted privileges. Do not choose **Interactive session** or **System** accounts unless absolutely necessary.
   Do note that even if you have prevented scripts from being run in your security policies, SES Evolution will assume that your internal custom scripts are trustworthy and allow them to be executed.
   Login to a domain
   Enter the name of the domain and indicate whether the agent needs to be connected for the test to be validated. A value of *Not connected* indicates that:

   • The agent is not linked to the domain in question,

   • If the agent is linked to the domain, that it is not connected to the domain network.

   Status of a network interface
   Click on ••• to add a network interface, and specify its name, unique ID (GUID) or description. Indicate as well what its status must be for the test to be validated: **Connected** or **Offline or idle**.
   The Name and Description of an interface can be seen in the Windows Network & Internet panel.
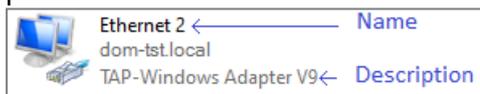
   Under Windows 10, to obtain all the information concerning an interface, including its GUID, run the following PowerShell command:
   ```
   Get-NetAdapter | Select Name, InterfaceName, InterfaceGUID,
   InterfaceDescription, Status
   ```

4. Add other tests if necessary, and click on **OK**. The sequence of the tests does not matter because ALL tests must be validated before the condition can be met.

5. Add other conditions if necessary. As soon as one condition is met, the corresponding policy will apply.
   Conditions apply in the order of their appearance.

6. If you want to run a custom script every time the conditional policy is applied, click on **Add a task**. When the script is added, specify its path, arguments and where to run it.
   It is best to use **Local service** as this is an account with restricted privileges. Do not choose **Interactive session** or **System** accounts unless absolutely necessary.

7. Under **Triggers**, select one or several events that will trigger the verification of conditions:

- Enable **Every** to check conditions at the regular interval that you specify.
- Enable **Network event** to check conditions if the workstation network interface changes, e.g., if it is connected to a Wi-Fi network, if it is a laptop plugged into a docking station, etc.

8. Click on **OK**. A summary of the conditions will appear in the **Policies** tab of the agent group.

9. Arrange the conditions in the sequence of your choice using the arrows on the left. The sequence of conditional policies is important.

✎ EXAMPLE 1
Quarantining a workstation if its health indicators are unsatisfactory.

In this example, every 10 minutes, a script will run on the agents and check their health status. If an agent's results are unsatisfactory, the *Quarantine* policy will be applied to the agent and a second repair script will run. A quarantine policy isolates an agent by blocking, for example, its communications over the network and all removable devices, except those used by administrators.

NAME	HealthCheck	✕

Policy: Quarantine

Conditions
At least one condition must be met
➕ Add a condition

HealthCheckScript	✕

Tasks
Tasks are optional
➕ Add a task

repair.ps1	✕

Triggers
Enable the triggers that make it possible to check the above conditions
Every 10 Minutes
Network event
✓ OK   ✕ CANCEL
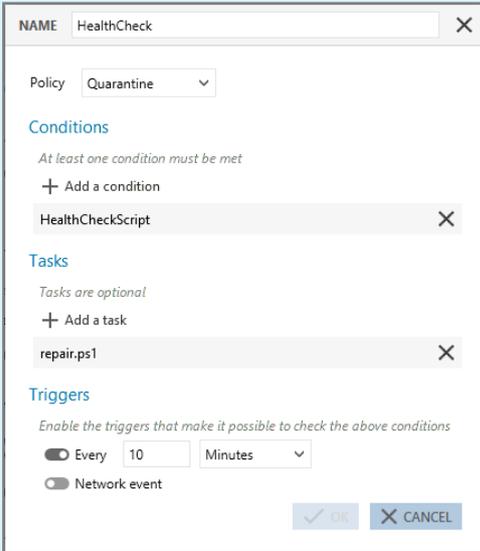
✎ EXAMPLE 2
Applying a specific policy for laptop computers.

In this example, every time a network event occurs on a workstation, SES Evolution will launch all the tests defined for this condition:

- The workstation is not connected to its domain network,
- The agent handler cannot be reached.

If the results of the tests are positive, the *Mobility* policy will be applied.

7.2.2 Enabling Windows shadow copies

SES Evolution's anti-ransomware protection mode keeps track of when files are modified and encrypted and blocks the process behind these operations if they are deemed malicious. Some files may nonetheless be encrypted before the process is effectively blocked.

If you enable anti-ransomware protection, you are strongly advised to enable the daily creation of shadow copies via SES Evolution. This feature, based on the Windows VSS service, will allow you to quickly restore the few lost files.

> ⚠ **WARNING:**
> Activating shadow copies cannot replace regular backups. You must have a dedicated parallel backup solution.

Requirements

You must meet the following Windows requirements in order to enable shadow copies in SES Evolution:

- Allow the creation of shadow copies for all NTFS volumes on all workstations protected by an SES Evolution agent.

- Reserve disk space for shadow copies on all local NTFS volumes on workstations protected by an SES Evolution agent.
  Use the Windows command `vssadmin resize shadowstorage` to set the size of this space.

For more information, refer to **Microsoft documentation**

> 📝 **EXAMPLE**
> Run the command:
> `vssadmin resize shadowstorage /For=C: /On=C: /MaxSize=15%`
> to reserve 15% of the space on the C:\ volume to store shadow copies on the C:\ volume.

Enabling shadow copies

1. In an agent group's **Policies** tab, go to **Daily shadow copies**.

2. Select the **Enable daily shadow copies** option.
   Every 24 hours, SES Evolution will make a shadow copy of local drives on the workstation running on an NTFS file system. Only the last five copies will be kept.

For more information on anti-ransomware protection and the process of restoring encrypted files, refer to:

- Configuring threat protection
- Managing ransomware attacks

7.2.3 Detecting and configuring the trust level on devices

SES Evolution monitors USB keys and other USB storage devices. Whenever a USB storage device is connected to an SES Evolution agent, it can be detected and displayed in the **Security > Devices** panel in the administration console depending on the options enabled. In this panel, a trust level can be manually assigned to these devices. For further information, refer to the section Changing the trust level of a USB device.

Some actions can also be applied automatically to all USB devices connected to agents in a group.

1. In an agent group's **Policies** tab, go to the **Device identification** section.

2. Enable **Allow device identification** if you wish SES Evolution to detect every USB device connected to an agent in the group and to automatically assign it the trust level 0 or 1. For the meaning of levels 0 and 1, see Changing the trust level of a USB device.

3. Go to the **Trusted devices** section.

4. Enable the **Trust empty devices** option if you want SES Evolution to detect every USB device connected to an agent in the group and automatically assign trust level 2 to every empty USB device.

5. Enable the option **Automatically scan devices** to automatically assign trust level 2 to every USB device connected to an agent in the group.
   When this option is enabled, the antivirus module(s) installed on the workstation will scan the key when it is plugged in, and neutralize potentially malicious files. If the antivirus is able to scan all the files, the device will be considered trustworthy. However, if some files cannot be accessed, the device will not be granted trust level 2, but will keep its current level.

For more information, please refer to the section Managing USB storage devices.

7.2.4 Creating scheduled tasks

Scheduled tasks make it possible to automatically run scripts on agents at regular intervals and/or when a network event occurs.

1. In the **Scheduled tasks** tab of an agent group, go to the section **Scheduled tasks** and click on **Add a scheduled task**.

2. Enter a name for the task in the **Run custom script** window.

a. To the right of the **Script** field, click on + to add the script to run.

b. In the **Arguments** field, specify the arguments to add when the script is run.

c. In the **Run in** list, choose **Local service** because this is an account with restricted privileges. Do not choose **Interactive session** or **System** accounts unless absolutely necessary.

4. Under **Triggers**, select one or several events that will trigger the execution of the script:

   - Enable **Every** to launch the script at the regular interval that you specify.
   - Enable **Network event** to run the script if the workstation network interface changes, e.g., if it is connected to a Wi-Fi network, if it is a laptop plugged into a docking station, etc.

5. Click on **Validate**.

All scripts that were declared in SES Evolution appear in the **Script** list. Select an existing script and click on [⊙] to view it or [⬆] to import a new version of the script.

7.2.5 Scheduling Yara scans

Scheduled scans make it possible to automatically run Yara scans on user workstations at regular intervals. For further information, refer to the section Running Yara scans.

To schedule scans, you must first create analysis units. For more information, refer to the section Creating Yara analysis units.

1. In the **Scheduled tasks** tab of the selected agent group, go to the section **Scheduled scans** and click on **Schedule a scan > Schedule a Yara scan**.

2. Enter a name for the scan in the **Schedule a scan** window.

3. Click on **Add analysis units** and select the analysis units that you want to include in your Yara scan. Click on **Next**.

4. Click on Log settings to determine the severity and destination of the SES Evolution logs generated during the Yara scan.

5. In **File scan parameters**, select **Default scan** to run a recursive scan on the folder\\.\EsaRoots\SystemDrive and exclude the folders \\.\EsaRoots\SystemRoot, \\.\EsaRoots\ProgramFiles and \\.\EsaRoots\ProgramFilesX86. Otherwise, select **Custom scan**:

   - **Analyze the image file of running processes**: checks whether the *.exe* file in the processes contains the Yara pattern you are looking for. This option also allows you to shut down any malicious processes identified on agents during the Yara scan, and/or exclude from the scan any processes run by Windows administrator and/or system accounts.
   - **File extensions:** Restricts scans to the indicated extensions.
   - **Included files and folders:** runs the scan on indicated files and folders with or without recursion.
   - **Excluded files and folders:** excludes from the scan indicated files and folders with or without recursion. Click on the + icon to add another path.

6. In the **Process scan parameters**, select **Default scan** to run a memory scan of all the processes being executed on the workstation, otherwise, select **Custom scan**:

   - **Shut down the process detected**: Stops dangerous processes identified during the Yara scan.
   - **Exclude processes run by:** Excludes from the analysis the processes that were run with the indicated integrity levels (administrator and/or system).
   - **Directory of excluded processes**: Excludes from the analysis the processes for which the executable files are located in the indicated folders. Click on the + icon to add another path.
     You can also export scan settings in JSON format and import them again for other tasks.

7. Fill in the information about the scheduled scan:
   - Period for which the scheduled scan will be active,
   - Frequency with which the scheduled scan will be run,
   - Time at which the scan starts. If the agent is not running at the indicated time, the scan will be launched as soon as the agent is restarted.

8. You can import all the settings of a scheduled scan that was exported earlier in JSON format.

9. Click on OK.

10. To deploy the scheduled scan on all agents in the group so that they apply it, go to the Security > Deployment menu and click on Deploy.

11. Read the agent's logs to ensure that the scans were run. You can also refer to the section on Looking up Yara scan usage.

7.2.6 Scheduling IoC scans

Scheduled scans make it possible to automatically run IoC on user workstations at regular intervals. For further information, refer to Searching for indicators of compromise.

To schedule scans, you must first create analysis units. For more information, refer to the section Creating IoC analysis units.

1. In the Scheduled tasks tab of the selected agent group, go to the section Scheduled scans and click on Schedule a scan > Schedule an IoC scan.

2. Enter a name for the scan in the Schedule a scan window.

3. Click on Add analysis units and select the analysis units that you want to include in your IoC scan. Click on Next.

4. Click on Log settings to determine the severity and destination of the SES Evolution logs generated during the IoC scan.
   The sections displayed below depend on the type of indicators in the analysis units selected in the previous step.

5. For Text indicators, you can disable the IoC scan in files, processes or event logs by unselecting the Text search checkboxes.

6. In File scan parameters, select Default scan to run a recursive scan on the folder\\.\EsaRoots\SystemDrive and exclude the folders \\.\EsaRoots\SystemRoot, \\.\EsaRoots\ProgramFiles and \\.\EsaRoots\ProgramFilesX86. Otherwise, select Custom scan:
   - Analyze the image file of running processes: checks whether the .exe file in the processes contains the indicators you are looking for. This option also allows you to shut down any malicious processes identified on agents during the IoC scan, and/or exclude from the scan any processes run by Windows administrator and/or system accounts.
   - File extensions: Restricts scans to the indicated extensions.
   - Included files and folders: runs the scan on indicated files and folders with or without recursion.
   - Excluded files and folders: excludes from the scan indicated files and folders with or without recursion. Click on the + icon to add another path.

7. In the **Process scan parameters**, select **Default scan** to run a memory scan of all the processes being executed on the workstation, otherwise, select **Custom scan**:

   - **Shut down the process detected**: Stops dangerous processes identified during the IoC scan.

   - **Exclude processes run by:** Excludes from the analysis the processes that were run with the indicated integrity levels (administrator and/or system).

   - **Directory of excluded processes**: Excludes from the analysis the processes for which the executable files are located in the indicated folders. Click on the + icon to add another path.

8. In the **Event logs** section, select the types of logs to scan and from which date.

9. In the **DNS request parameter** section, indicate the date from which you want to analyze DNS requests.

10. Fill in the information about the scheduled scan:

    - Period for which the scheduled scan will be active,

    - Frequency with which the scheduled scan will be run,

    - Time at which the scan starts. If the agent is not running at the indicated time, the scan will be launched as soon as the agent is restarted.

11. You can import all the settings of a scheduled scan that was exported earlier in JSON format.

12. Click on **OK**.

13. To deploy the scheduled scan on all agents in the group so that they apply it, go to the **Security > Deployment** menu and click on **Deploy**.

14. Read the agent's logs to ensure that the scans were run. You can also refer to the section on **Looking up IoC scan usage**.

7.2.7 Understanding self-protection on agents and performing maintenance operations

SES Evolution agents are equipped with a self-protection mechanism implemented by a set of rules that are transparent for administrators and users. With these rules, you can:

- guarantee that the security policies applied by administrators do not hinder the proper operation of agents (but does not prevent policies from hindering the operation of workstations when there are wrongly configured rules),

- protect agents from external attacks or malicious users who may attempt to disable or uninstall the agents.

However, to perform maintenance operations on agents in a group, they must first be switched to Maintenance mode to disable the self-protection system. To do so, you must allow Maintenance mode to be used in the group's configuration.

Enabling Maintenance mode

Administration privileges are required to enable Maintenance mode.

All maintenance operations performed will be logged while Maintenance mode is enabled.

The agent's automatic updates will also be suspended when Maintenance mode is enabled. They will be applied automatically when Maintenance mode ends. You can also apply forced updates. For more information, see the section **Forcing an update on agents**.

> ⊘ **CAUTION**
> When Maintenance mode is enabled, the agent continues to protect the workstation because the

security policy stays enabled. However, this mode must be used with caution and by trustworthy users.

1. In an agent group's **Settings** tab, go to the **Maintenance** section.
2. Enable the parameter **Allow Maintenance mode**.
3. Deploy the configuration in the environment to apply the new configuration.

The user must enable Maintenance mode in the agent's interface on their side, in the **Helpdesk** tab of the **Help and support** panel ⌸. For further information, refer to the section Configuring preferences on the agent.

You can also enable Maintenance mode individually on the workstation concerned using challenges. Administration privileges are not required. For further information, refer to the section Resolving issues with challenges.

Disabling Maintenance mode

When maintenance operations are completed, remember to end Maintenance mode by clicking on **Disable** in the agent's interface to restore self-protection and security. The integrity of the agent's resources will then be checked. If anomalies are detected, the agent will launch repairs. The user may then be asked to restart the workstation.

You can also enable and disable Maintenance mode via a script, by launching EsGui ([…]\Stormshield\SES Evolution\Agent\Bin\Gui) with the command line options `/EnterMaintenanceMode` and `/LeaveMaintenanceMode`.

Administration privileges are not required to disable Maintenance mode.

7.2.8 Allowing administrators to uninstall agents

The only way to uninstall an SES Evolution agent from a user workstation by default is via a challenge. For further information, refer to the section Uninstalling an agent.

However, an agent group can be configured to allow a workstation administrator to uninstall the SES Evolution agent without a challenge.

1. In an agent group's **Settings** tab, go to the **Uninstallation** section.
2. Enable the **Allow agent uninstallation** parameter and click on **Save**.
3. Deploy the configuration in the environment to apply the new configuration.

7.2.9 Collecting diagnostic data

By default, the only solution to collect diagnostic data on user workstations that cause issues is the challenge. For further information, refer to the section Running a diagnostic.

However, an agent group can be configured to allow a workstation administrator to run a diagnostic without a challenge.

1. In an agent group's **Settings** tab, go to the **Collecting diagnostic data** section.
2. Enable the **Allow collection of diagnostic data** setting and click on **Save**.
3. Deploy the configuration in the environment to apply the new configuration.

For further information on diagnostics, refer to the section Troubleshooting issues.

7.2.10 Choosing agent update settings

1. In an agent group's **Settings** tab, go to the **Version** section.
2. In **Version**, choose the version of the agent to apply to this agent group.
3. Enable the option **Allow downgrading to older version** to allow updates to lower versions of the agent.
   This option is particularly useful if you find an operating problem with one version of the agent. It allows you to return to a previous version on which the problem does not appear.
4. Disable the **Apply software updates automatically** option to prevent the agent handlers from applying the updates during a new deployment over the pool of agents. In this case, only configuration or security policy changes are deployed, when compatible with the current version of the handlers.

For more information, see Updating agents, particularly to know the other agent updating modes if you have disabled the second option.

7.2.11 Choosing the features to enable on agents

To avoid incompatibility issues or duplicates with other installed programs, some features on SES Evolution may need to be disabled.

1. In an agent group's **Settings** tab, go to the **Active features** section.
2. Unselect the features that you want to disable.
   After the new configuration is applied, a message on the agents' dashboard will indicate that these agents need to be restarted.

The agent interface displays the list of features, and their status, in the **Protections** tab of the **Help and Support** panel. For more information, see Getting help on the agent.

7.2.12 Choosing the agent handler groups assigned to agents

You can choose the agent handlers to which agents in a group must connect to send their information and retrieve various updates. If your infrastructure is spread out over several physical sites, it may be useful to distribute agent groups to the closest agent handlers.

Agents that are not associated with any agent handlers are known as standalone agents. All of their updates must be performed manually by generating an installer and running it on agents, in the same way as during their initial deployment. For more information, see Installing agents on workstations.

1. In an agent group's **Settings** tab, go to the **Agent handlers** section.
2. In **Default agent handler group**, add the agent handler group(s) to which the agents of this agent group must connect.
3. In **Backup agent handler group**, add the agent handler group(s) to which the agents can connect if default groups fail.

7.2.13 Showing Technical support information on agents

The information shown in the tab **Help and support > Contact** can be customized in the agent's interface.

1. In an agent group's **Settings** tab, go to the **Help and support** section.

2. Enter the description that you want to display in the header of the **Contact** tab in the **Help and Support** panel of the agent's interface, e.g., "*If you encounter issues with SES Evolution, feel free to get in touch with the IT department*".

3. Enter the **Email address**, **Telephone** number and **Website** of the department that manages technical support for SES Evolution.

7.2.14 Monitoring agents in real time

In the table of the **Agents** tab, you can distinguish connected agents from offline agents. Offline agents are grayed out

Agents are considered offline if they have not connected to an agent handler for the duration defined in the agent group configuration.

You can set the frequency with which agents connect to the agent handler to update their status. You can also customize the duration as well as how long agents can stay offline before they will be automatically deleted from the database. For further information, refer to the section Removing offline agents automatically.

To set the connection frequency and the various durations:

1. In an agent group's **Status and logs** tab, go to the **Agent real time monitoring** section.

2. Choose the frequency of the **Agent status update** in seconds. The agent connects automatically to the agent handler by default every 60 seconds to:
   - Send information about its status to refresh the agent group panel,
   - Retrieve new configurations, policies or updates if there are any.

   You can also manually force a connection to the agent handler and log sending by clicking on **Check for updates** in **Protection status** in the agent's interface.

3. Set the value of the **Disconnection after** setting. By default, agents are considered offline if they have not connected to their agent handler for seven consecutive days.

4. Set the value of the **Automatic deletion after** parameter. Agents are deleted by default after 30 consecutive days of staying offline.

7.2.15 Sending logs generated by agents

1. In the **Status and logs** tab of an agent group, go to the **Logs** section.
2. Choose the severity level above which logs will be sent to the following destinations:
   - **Show the agent** in the **Help and Support** panel, **Events** tab of the agent interface,
   - **Show on console** in the **Environment > Agent logs** panel of the administration console, i.e. stored in the log database.

   For example, if you choose *Informational* for the agent, all logs can be viewed in the agent's interface, except for *Debug* logs.
   Emergency and *Alert* logs are always sent to all destinations. Logs that are not sent can never be read.
   Note that only *Alert* and *Emergency* level logs that have led to a block are visible in the agent interface for a non-administrator user of his machine.
   If you are validating new software, a new workstation etc., temporarily transmit the *Information* level logs. In the event of maintenance or troubleshooting, *Debug* logs will also be useful.
   For more information on log severity levels, refer to the section Monitoring SES Evolution agent activity.
   To refine this global action, you can define the logs to send for each security rule. For further information, refer to the section Configuring log management.
   To configure how logs are sent to Syslog servers, see Creating groups of agent handlers.
3. In the **Log transmission frequency** section, choose the maximum frequency (in seconds) with which the agent's logs will be sent to the agent handler:
   - **Urgent logs** correspond to *Emergency* and *Alert* logs.
   - **Standard logs** group all other levels.

   This parameter allows you to manage bandwidth use. Urgent logs are sent every 30 seconds by default and standard logs are sent every hour (3600 seconds).
4. Logs displayed on an agent are deleted from the disk by default based on the following criteria:
   - When logs exceed 500 MB. In this case, the oldest logs will be deleted until they occupy less than 500 MB.
   - When logs are more than 30 days old.
     This duration can be modified in the **Keep logs of less than** field. If this option is fully disabled, only the file size criterion will apply.
5. If required, choose **Send agent self-protection logs** to the agent handlers. These are logs collected from the various mechanisms that protect components essential to the integrity of the agent. When this parameter is disabled, self-protection logs will remain available on agents.

7.2.16 Configuring context details generated by agents

Context details are all the logs that the agent produced in an attack perimeter, including those that do not usually appear in the administration console. For example, even logs that remained local on the agent or that were sent to a syslog server are shown in the context details. For further information, refer to the section Analyzing contexts to understand attacks.

You can configure the size of such contexts, the maximum age of their logs and how they are sent to the agent handler.

1. In an agent group's **Status and logs** tab, go to the **Context** section.

2. Define the **Size limit** of a context, which is 500 KB by default. This is the estimated size of data going through the network. If network connections are restricted between agents and the agent handler, reduce this size. Conversely, if you added highly verbose sets of audit rules, increase this size to ensure that you retrieve enough useful logs.

3. Define the **Oldest logs**. The default value is 10 minutes because most attacks happen quickly, but you can adjust it according to your preferences.

4. Choose how **Context detail reporting** takes place from the agent to the agent handler. Reporting can be:

   - **Immediate**: context logs are sent to the agent handler at the same time as the alert, and can be seen immediately in the administration console.

   - **Postponed**: context logs are sent to the agent handler at a **Frequency** that can be defined, the default value being every hour. If you analyze attacks only once daily, increase this frequency to every two or three hours to avoid network congestion.

   - **On demand**: context logs will not be sent to the agent handler automatically. You can download all this data manually when you intend to analyze an attack. For further information, refer to the section Analyzing contexts to understand attacks.

5. Save your changes.

7.3 Installing agents on workstations

As soon as you have configured your agent groups, you must install agents on the workstations that you want to protect.

An SES Evolution agent can be installed on all types of hosts with compatible operating systems: servers or workstations, including domain controllers or machines that host one or several SES Evolution components (such as agent handlers, backends, etc.)

This installation is a two-step process. First, generate an installer that contains the whole configuration dedicated to the agent group. Next, deploy the agent on each workstation you want to assign to this group. Once it is installed, the agent will retrieve a unique identity the first time it connects to the agent handler. It will then appear in the panel of the corresponding agent group in the administration console. The whole configuration of the agent group will be applied to it, especially security policies.

If you have installed SES Evolution on a master, you also need to change the ID of the agents on which you are deploying it. For further information, refer to Installing the agent on workstations from a master in the Administration Guide.

> 🛈 **NOTE**
> The folder directed by %TEMP% and %TMP% must exist and be accessible in write mode during the agent installation phase and during the agent update.

Regardless of the agent deployment method chosen, once installation is complete, ensure that the agents connect at least once to an agent handler.
An agent that has never logged on to an agent handler performs its workstation protection functions, but the following features are not available:

- log generation. You will not be able to monitor its activity or export unsynchronized logs from the **Help and Support** > **Assistance** menu of the agent interface.

- challenge request. You will not be able to perform a challenge if there is a problem. For further information on challenges, refer to the section Resolving issues with challenges.

7.3.1 System requirements for agents

To install and use Stormshield Endpoint Security Evolution version 2.7.2 on Microsoft Windows, agents must meet at least the following requirements:

Operating systems	Refer to the Product life cycle guide to find out more on compatibility with operating system versions.
Processors for physical machines	64-bit processors with minimum 2 GHz Intel Pentium 4 or equivalent. Itanium processors are not supported.
Processors for virtual machines	At least one virtual socket and one 1 GHz core per socket. Stormshield recommends one virtual socket and two 2 GHz cores per socket.
Physical memory	At least 1 GB. Or more if the operating system requires it. Stormshield recommends 2 GB.
Disk space	- At least 100 MB for installation, - At least 200 MB for data storage. These are the disk space requirements for the NTFS file system. More space will be needed for updates and logs.
Network configuration	- Outgoing communication: - Agent with version 2.7.1 or higher - port HTTPS 443 - Agent with version lower than 2.7.1 - port TCP 17000
Network bandwidth	At least 12 Kbit/s. Lower bandwidth may prevent the agent and agent handler from exchanging data.
Software	Framework .NET 4.6.2 or higher.
Display	At least 1024X768.

Enabling Windows restore points

The SES Evolution agent installer creates a Windows restore point just before copying files on the disk. So if there are any compatibility issues with another program, this will make it possible to revert to the state of the system as it was before SES Evolution as installed. A restore point will also be created when the agent is updated.

In order for the restore point to be created, the feature must be enabled in the System > System protection panel in Windows. For further information on restoration, refer to Windows documentation.

Disabling safe mode for standard users

Safe mode can be used to troubleshoot problems that prevent a workstation from being used when started normally. By default, the Windows configuration allows all users to start in this mode.

However, in safe mode, the SES Evolution agent self-protection is disabled. You must therefore allow only administrators to use this mode.

To disable safe mode for non-administrator users, set the *SafeModeBlockNonAdmins* value of the *HKLM\Software\Microsoft\Windows\CurrentVersion\Policies\System* key to "1" in the Windows registry.

7.3.2 Generating an installer for agents

The **Agent groups - Modify** privilege is required to generate an installer for agents.

1. Select the **Environment > Agents** menu.
2. Ensure that you have configured the agent group with your preferences and deployed the environment. For further information, refer to the section Creating agent groups.
3. From the panel on the left, select the agent groups that you want to apply to the workstations.
4. In the **Agents** tab, click on **Installer > Generate xx-bit installer**.
5. Save the installation file *AgentSetup_xxx.exe* at the location of your choice.

7.3.3 Deploying the agent to each standard workstation via GPO

Once the installer has been generated, you can deploy this file to the workstations via GPO. The GPO-based procedure described below uses the PowerShell script *SesAgentDeploymentScript.ps1* provided by Stormshield, and launches a default installation in silent mode.

1. In your MyStormshield client area, select **Downloads > Stormshield Endpoint Security > Evolution > Resources** and click on the *SES Agent deployment script* link. The script requires PowerShell in version 5 and higher.
2. On the domain controller, open the Group Policy Management console (*gpmc.msc*).
3. Right-click on the organizational unit in which you want to deploy the SES Evolution agent, then select **Create a GPO in this domain, and link it here**.
4. In the **New GPO** window, enter a name for the GPO, e.g., *SES EVOLUTION Deployment*.
5. Right-click on the new GPO, then select **Edit**.
   The Group Policy Management editor then opens.
6. Select **Computer configuration > Policies > Windows settings > Scripts (Startup/Shutdown)**, and double-click on **Startup**.
7. In the **Startup properties** window, click on the **PowerShell scripts** tab, then on **Show files** and paste the following files:
   - The *AgentSetup_x64.exe files*,
   - The *SesAgentDeploymentScript.ps1* script.
8. Click on **Add**, then on **Browse**
9. Select the script, click on **Open**, then on **OK**.
10. In the **Startup properties** window, click on **Apply**, then **OK**.
11. In the Group Policy Management console, select the GPO that was just created.
12. In the **Scope** tab, check the following items:
    - The organizational unit in the **Links** section,
    - The target user groups listed in the **Security filtering** security.
13. Right-click on the OU, then select **Group Policy Update**.
    The SES Evolution agent will automatically install in silent mode the next time the workstations start.
    You can refer to the logs regarding the installation via GPO in
    C:\Windows\Temp\InstallSESLogGPO.

14. As soon as the agent is installed, the icon ⬛ appears in the Windows status bar, indicating that the installation is not complete.

15. Restart the workstation. The icon ⬛ indicates that the agent is now fully functional.

7.3.4 Deploying the agent on each standard workstation via MECM (formerly SCCM)

Once the installer has been generated, you can deploy it on the workstations via Microsoft Endpoint Configuration Manager, which replaces SCCM.

> **ℹ NOTE**
> To deploy the agent via MECM, SES Evolution must be in at least version 2.3 and built-in security policies must be in at least version 2210a.
> Go to your MyStormshield client area to download the most recent versions of SES Evolution and policies. You can also download the latest policies from the update server. For more information, see Downloading Stormshield updates.

In your MECM environment, we recommend that you:

- Have at least one shared folder that can be used by hosts in the pool connected to MECM,
- Split up the list of hosts in the pool into **Collections**. You can divide SES Evolution agents into agent groups, for example.

The following procedure was tested on version 2207 of MECM.

To deploy the agent via MECM, follow the four steps below:

1	Creating an installation package	1. Place the installation files *AgentSetup_x64.exe* in the shared folder usable by machines in the pool connected to MECM.
		2. Open the **Configuration Manager** console.
		3. In **Software library** > **Overview** > **Application management**, click on **Packages**.
		4. Create a package and a program. Name the package, e.g., *InstallSES Evolution agent*.
		5. Select the standard program type.
		6. Name the program, e.g., *SES Evolutionagent on Windows x64*.
		7. Enter command line `AgentSetup_x64.exe /s`. Other options are available, see the list of options below the table.
		8. Go to the shared folder containing the installation file, and select it as the start folder.
		9. Select **Run with administrative rights** as the run mode.
		10. In **Platform requirements**, select the operating system on which the agent will be deployed.
2	Creating programs to install via the package	In the package created, create as many programs as necessary, for each agent group for example. To create a new program:
		1. Right-click on the package and select **Create program**.
		2. Select the parameter as indicated in step 1.

3	Deploying programs on workstations	1. Right-click on each program created in the package, and select **Deploy**.
		2. Select the **Collection** in which you intend to deploy the agent.
		3. Specify a deployment **Scheduling**.
		4. In **Notification settings**, select **Allow users to run the program independently of assignments**.
		5. In **Deployment options**, select **Run program from distribution point**.
4	Monitoring and finalizing the deployment	1. In **Monitoring** > **Overview**, click on **Deployments**.
		2. Select an ongoing deployment to view its progress.
		3. You can select **Run summarization** to force the workstations to synchronize with MECM.
		4. As soon as the agent is installed on the workstations, the icon appears in the Windows status bar, indicating that the installation is not complete.
		5. Restart the workstations. The icon indicates that the agent is now fully functional.

You can also add the following options to the `AgentSetup_x64.exe` command:

`/silent` or `/s`	To make the installation transparent for the user of the workstation
`/installdir`	To copy the agent's installation files (binary and resource files) into a folder other than *%SYSTEMDRIVE%\Program Files*. This path must be different from the one for the agent's data files.
`/datadir`	To copy the agent's data files (logs, policies, scripts, etc.) into a folder other than *%SYSTEMDRIVE%\ProgramData*. This path must be different from the one for the agent's installation files.
`/log <path>`	To specify the path of the agent's installation log file.
`/newagentid`	To delete data regarding the agent's communication with the agent handler: unique ID, certificates used internally, as well as the ID and private data used in challenges. The agent retrieves new data the next time it connects to the agent handler. It would be helpful to assign new communication data if the agent is on a duplicated virtual machine, or if you are installing the agent on a workstation created from a master.

7.3.5 Installing the agent on workstations from a master

1. Install an SES Evolution agent on a master by following the procedure for the installation of a standard agent.

2. On the master, delete the agent's ID by following one of the methods below. Agent handlers must not be contactable by the agent during this operation, otherwise the agent will immediately obtain new communication data.

- Delete the registry value of the agent's ID (value: *AgentGuid*) located in: *HKEY_LOCAL_MACHINE\SOFTWARE\Stormshield\SES Evolution*. A new ID will be generated the next time the agent connects to the agent handler.
  - or -

- Run the agent installer *AgentSetup_x64.exe* or the agent component *Agent\bin\Gui\EsSetup.exe* in command mode with the */newagentid* option. This command assigns a new ID to the agent without the need to reinstall it.

After deploying the master to a workstation, ensure that the SES Evolution agent can contact an agent handler so that a new ID can be assigned to it.

7.3.6 Using agents on Microsoft Windows Server Core operating systems

You can install the SES Evolution agent on Windows Server Core 2012 R2, 2016, 2019 and 2022 operating systems.

These operating systems have a reduced graphical interface. The agent's interface does not start automatically when a user session is opened (icon in the task bar on a 'standard' operating system). To display the agent's GUI:

- Use the command `EsGui.exe`.

Likewise, if requests for user confirmation are configured in a security rule, the agent will not open any window, automatically assuming that the answer to the confirmation is "no". There is no way for the user to reply with a "yes".

7.4 Viewing agents in the console

The administration console allows you to track the status of agents in real time on all workstations. Agents can be classified by various criteria: operating system, domain, SES Evolution version, etc.

Agents can also be filtered, moved from one group to another and exported to a CSV file.

Users need to hold the **Agent groups - Display** privilege to view this panel.

7.4.1 Displaying the agent list

1. In the **Environment > Agents** menu, select **All agents** to view all agents regardless of their group.
   - or -
   Select an agent group from the left panel, then click on **Agents**. Every agent deployed via the agent group installer connects to the agent handler and appears in the table with the following information:

   - **Computer**: name of the workstation on which the SES Evolution agent is installed,

   - **IP address**: Main IP address if the computer has several network cards,

   - **Version**: version number of the SES Evolution agent,

   - **Operating system**: Workstation operating system version,

   - **Host type**: desktop PC, laptop, server, virtual machine, or unknown,

- **Policy**: name of the SES Evolution security policy applied to the workstation,
- **Last connection**: date of the SES Evolution agent's last connection to the agent handler,
- **Domain**: name of the Windows domain to which the workstation belongs,
- **User**: name of the account that last connected to the SES Evolution server from this workstation,
- **Group**: Name of the agent group to which the agent belongs,
- **Mode**: operating mode of the SES Evolution agent: normal, stopped or maintenance. Stopped mode means that SES Evolution no longer protects the workstation. For further information on Maintenance mode, refer to the section Enabling Maintenance mode.
- **Pinned**: the ⬛ icon means that the agent will remain in its agent group regardless of Active Directory assignment rules. If the column is empty, the agent will comply with Active Directory rules and it can be moved automatically from one group to another if its Active Directory criteria change. For further information, refer to the section Automatically assigning agents to agent groups.

2. Click on a column title to sort the list of agents by this criterion. For example, click on **Group** to sort agents by their agent group.

7.4.2 Filtering the list of agents

1. In the **Filters** section of the **Agents** tab, enable filters to customize your list of agents. Every column corresponds to a type of filter and contains several values. Click on these values to enable the corresponding filter.
   The list of agents will be refreshed according to the active filters applied.
2. You can go back to the full list of agents at any time by clicking on **Reset filters**.

To filter by a computer's name, its GUID, date of last connection or by user, enter a character string in the search field at the top on the right.

7.4.3 Moving agents from one group to another

To move a selection of agents:

1. In the list of agents, select the agents that you want to move.
2. Click **Move selected agents to > Desired group name**. Agent names are displayed in blue italics in both the source and target groups to indicate that agents are being moved.

To move all displayed agents:

- Click on **Move displayed agents to > Desired group name**. Agent names are displayed in blue italics in both the source and target groups to indicate that agents are being moved.

After moving agents, you must deploy to apply the configuration and security policies of the new group to the agents:

- Select the **Security** > **Deployment** menu, and click on **Deploy**.
  Agent names are displayed in black again. They are removed from the group and now belong to the group to which they have been moved.

If the agent has been placed in its group by an Active Directory assignment rule and you move it manually, it will be pinned to its new group. The Active Directory assignment rule will no longer apply.
If the agent was pinned in its original group, it will be pinned in its destination group. For further

information on Active Directory assignment rules and pinning, refer to the section Automatically assigning agents to agent groups.

7.4.4 Exporting a list of agents

Information on agents can be exported to a CSV file so that it can be read and processed in a spreadsheet.

1. In the list of agents, select the agents that you want to export.
2. Right-click and choose **Export selected agents**, then select the desired separator, i.e., comma, semicolon or tab. A file *ExportedAgents.csv* is created by default on the desktop. Change its name and destination if necessary.
3. Open the *.csv* file with the tool of your choice.

To monitor activity on your agents, you can view their logs. For further information, refer to the section Viewing and managing agent logs in the administration console.

7.5 Automatically assigning agents to agent groups

Agents can be automatically assigned to an agent group based on the Active Directory groups or organizational units to which they belong.

If you are using this feature, an agent will automatically be assigned to an agent group based on the Active Directory criteria on the agent when the workstation starts up:

- If the agent's Active Directory group or organizational unit changes later, the agent will be moved to the corresponding agent group after the workstation is restarted,
- If only assignment rules have been changed and deployed from the administration console, the agent will automatically be moved to the corresponding group without restarting the workstation.

To automatically assign agents, you must create assignment rules based either on Active Directory groups or organizational units. Verifications will be based on the Active Directory criteria of the host, not of the connected user.

If you want agents to continue being in an agent group regardless of their Active Directory criteria, you can pin them manually to this group.

The **Agent groups - Modify** privilege is required to create assignment rules.

7.5.1 Creating an agent group assignment rule

1. Select **All agents** in the **Environment > Agents** menu, then the **Assignment rules** tab.
2. Click on **Edit** at the top right.
3. Click on **Add rule based on AD group** or **Add rule based on OU**.
   A new row appears.
4. Enter a **Description** that would make it easy to recognize this rule.
5. Click on [...] and select the desired group or OU in the window that appears.
   You can also manually enter the group or organizational unit using LDAP syntax, e.g., *OU=Paris,DC=Grey,DC=local*.
6. In the **Assign to agent group** list, select the agent group to which the workstations of this group or OU will belong.

7. Create other rules if necessary.

8. Change the order of rules by scrolling over them to show the arrows on the left. If there are several rules that match an agent's AD criteria, the agent will be assigned to the agent group in the first matching rule.

9. Click on **Save**.

   The icon of the agent group in the left panel becomes 🔳 , showing that at least one Active Directory assignment rule affects the group.
   In the **Agents** tab of the agent groups concerned, the names of the assignment rules appear as links that provide direct access to the rules.

10. Select the **Security > Deployment** menu, and click on **Deploy**.
    An agent will be assigned to its group based on its AD create once the agent has retrieved the new configuration and sent back its AD criteria to the agent handler. The workstation may need to be restarted if the changes were made on the Active Directive controller during the user's session.

7.5.2 Pinning an agent to an agent group to ignore its Active Directory criteria

Manually pin agents to an agent group if you want them to keep their group regardless of their Active Directory criteria.

1. In the list of agents, select the agents that you want to pin.

2. Click on **Pin or unpin agents** > **Pin to group**.

   The 📌 icon is displayed in the **Pinned** column. The agent will continue to belong to this agent group no matter what, even if its Active Directory criteria change. It can only change groups if you move the agent manually or unpin it from the group.

3. Select the **Security > Deployment** menu, and click on **Deploy**.

7.5.3 Unpinning an agent from an agent group

Unpin an agent from an agent group if you want it to be assigned automatically to an agent group again based on its Active Directory criteria.

1. In the list of agents, select the agents that you want to unpin from the group.

2. Click on **Pin or unpin agents** > **Unpin from group**.

   The 📌 icon disappears from the **Pinned** column. The agent can now automatically change groups if an Active Directory assignment rule affects it.

3. Select the **Security > Deployment** menu, and click on **Deploy**.

7.6 Understanding the agent interface on workstations

The interface dedicated to agents displays information on the health status of each agent and analyzes event logs. This makes it possible to troubleshoot issues when they occur.

Support and diagnostic tools, as well as quarantine tracking, are also available.

- To open the SES Evolution agent interface, double-click on 🛡️ in the workstation taskbar.

7.6.1 Viewing the health status of an agent

The agent's **Protection status** dashboard displays the operation of the agent's four main protection modules according to the following color code:

- **Green**: all modules are running,
- **Orange**: A protection in a module is stopped or waiting for restart,
- **Red**: a module is not functioning,
- **Gray**: A protection in a module has been disabled in the agent group configuration.

Click on the module names or on the shield to access the protection status details. The **Protections** tab in the **Help and Support** panel appears. It shows the status of modules and their protections. This list of protections corresponds to the features you can enable or disable in the agent configuration in the administration console. For more information, see Choosing the features to enable on agents.

The middle area displays the last three *Alert* or *Emergency* level events, which resulted in a block, that were issued on the agent. Identical events are grouped together.

The lower part of the dashboard provides details about the configuration of SES Evolution:

- **Agent group**: name of the agent group to which this agent belongs,
- **Security policy**: name of the security policy applied to this agent,
- **Last policy update**: date on which this agent updated its security policy for the last time,
- **Last connection**: Date on which this agent last connected to the agent handler.

The agent connects automatically to the agent handler by default as frequently as configured. Click on **Check for updates** if you want agents to connect to the agent handler and perform the following operations:

- Send data about the status of the agent to the agent handler, including logs,
- Retrieve new configurations, policies or updates if there are any.

You can also force connection to the agent handler via a script, by running the EsGui program ([...]\StormshieldSES EvolutionAgentBin\Gui) with the command line option `/ForceConnection`.

7.6.2 Configuring preferences on the agent

1. Click on the ⚙ tab in the agent interface to open the **Preferences** panel.
2. Set the options according to your preferences. You can:
   - Choose the language of the agent interface,
   - Save the position and size of the agent interface window,
   - Show notifications in tool tips.

7.6.3 Getting help on the agent

- Click on the 🧰 tab in the agent interface to open the **Help and support** panel.

Viewing information about the agent

In the **Contact** tab, locate the following information:

- **Technical support contact**: details of the service to get in touch with if you encounter issues with the SES Evolution agent. This information appears only if you have chosen to do so in the configuration of the agent group. For further information, refer to the section Showing Technical support information on agents.
- **Information**: Details of the agent installed on this workstation.

Requesting assistance

In the **Helpdesk** tab, these three operations can be performed:

- Request a challenge from your administrator. For further information, refer to the section Resolving issues with challenges.
- Enable Maintenance mode if you need to disable self-protection on the agent to perform maintenance operations. Administrator privileges are required. For further information, refer to the section Understanding self-protection on agents and performing maintenance operations.

> ⚠️ **CAUTION**
> If Maintenance mode is enabled, it must be disabled as soon as maintenance operations are over by clicking on **Disable** in the agent interface. This will restore self-protection and security on the workstation. In addition, the agent's automatic updates are suspended when Maintenance mode is enabled.

- Export logs that have not yet been sent to the agent handler because the agent was disconnected from the network, for example.

Troubleshooting issues

To form a diagnosis, refer to the section Diagnosing issues on agents.

Viewing event logs

Agents' logs can be read in the **Events** tab of each agent's interface. For further information, refer to the section Viewing logs in the agents' interface.

Only *Alert* and *Emergency* level logs that have resulted in a block are displayed on a non-administrator user's agent.

An administrator user can view all log levels, blocking or non-blocking, according to the levels configured in the security rules to be sent to the agent. For further information, refer to the section Configuring log management.

Similar logs are grouped together in the agent interface of a non-administrator user.

Agent logs can also be read on the administration console and the syslog server, if you have configured one.

Monitoring quarantined files

The agent dashboard shows the number of quarantined files in the middle area. Click on the hint to display the **Quarantine** tab of the **Help and Support** panel.

The tab displays a real-time list of files that have been quarantined, with their name, location, quarantine date and the date on which they will be deleted.

Viewing protection status

In the **Protections** tab, you can view the list and status of SES Evolution agent protection modules. It corresponds to the agent configuration features you can enable or disable in the administration console. For more information, see Choosing the features to enable on agents.

Disabled modules or protections are displayed in gray or with a crossed-out round icon.

7.6.4 Using the EsGui command

The EsGui command makes it possible to interact with the interface of the SES Evolution agent. It can be used in scripts or in the shortcut to the interface on the Windows Desktop. The following are the options that this command offers:

`/silent or /s`	Hides the agent interface (only its icon) in the Windows notification area.
`/ShowPanel <panel>`	Shows a specific tab of the agent interface. The possible values are:

- `ProtectionStatus`: shows the **Protection status** dashboard,
- `Settings`: shows the **Preferences** panel,
- `Contact`: shows the **Contacts** tab of the Help and support panel,
- `Helpdesk`: shows the **Helpdesk** tab of the Help and support panel.
- `Troubleshooting`: shows the **Diagnosis** tab of the Help and support panel,
- `Logs`: shows the **Events** tab of the Help and support panel.

`/EnterMaintenanceMode`	Starts maintenance mode. The agent interface must be run with administration privileges, and maintenance mode must be allowed in the agent group settings. For further information, refer to the section Understanding self-protection on agents and performing maintenance operations.
`/LeaveMaintenanceMode`	Quits maintenance mode. The agent interface must be run with administration privileges.
`/GenerateDiagnostic <path.zip>` `/AcknowledgePersonalDataCollection` `/DiagnosticComment <comment>`	Generates a diagnostic package without tracing. Specify the destination folder of the package. The `/AcknowledgePersonalDataCollection` parameter is mandatory to allow the collection of personal data. The `/DiagnosticComment` parameter is optional, and makes it possible to add comments. For further information, refer to the section Troubleshooting issues.
`/StartDiagnosticWithTraces` `/AcknowledgePersonalDataCollection`	Starts diagnosing and tracing.
`/StopDiagnosticWithTraces <path.zip> /DiagnosticComment <comment>`	Stops tracing and finishes generating the diagnostic package. Specify the location to save the package.
`/CancelDiagnostic`	Cancels the diagnostic in progress.

`/GrantWebAccess`	Starts a period of temporary web access, if the policy allows it. For further information, refer to the section Accessing the web temporarily from the agent.
`/ExportLogs`	Exports the agent's logs. You can specify the destination folder as an option. The agent interface must be run with administration privileges. For further information, refer to the section Reading logs of offline agents.
`/ForceConnection`	Forces the agent's connection to the agent handler to send logs and information about the status of the agent, and to get new configurations, policies or updates. This option is the equivalent of clicking on the **Check for updates** button in the agent interface. For further information, refer to the section Viewing the health status of an agent.

7.7 Updating agents

There are two types of agent updates:

- Upgrade to a standard version of SES Evolution which distributes the new SES Evolution version to the agents.
- Upgrade to an LTSB version of SES Evolution which enables a patch version to be distributed to agents on workstations with operating systems not supported as standard. For further information, see the SES Evolution product life cycle document.

Once you have updated SES Evolution with the Installation Center or imported an LTSB patch version, you can then apply this version to one or more agent groups via the administration console. If some agents are not connected to the agent handler or if you do not require an automatic update, apply the new version to them manually.

An update should be applied to a group of test agents first, in order to test it. You can then apply it to your production groups.

To downgrade agents to an earlier software version of SES Evolution, ensure that the option **Allow downgrading to older version** is enabled in Choosing agent update settings.

The **Agent groups - Modify** privilege is required to update agents.

7.7.1 Importing an agent patch (LTSB version only)

This procedure only applies if you want to update an agent group with the LTSB version of SES Evolution.

1. In your MyStormshield client area, **Downloads** section, download the LTSB agent patch.
2. Select **All agents** in the **Environment > Agents** menu.
3. On the **Agents** tab, click on **Patch versions > Import an agent patch**.
4. Select the *.zip.p7* file, then click on **Open**.
5. Select the agent group to update to LTSB version.
6. On the **Settings > Version** tab, select version 2.5.7 (LTSB) and disable the **Automatically apply updates** option.
7. For manual updates, see Manually updating an agent and Force updating an agent.

7.7.2 Applying updates to agents that are connected to the agent handler

This procedure applies if you want the agent handlers to automatically update agents when a new deployment is performed. Otherwise, disable the **Automatically apply software updates** option in the menu shown below. See the two sections below for manual updates.

1. Select the **Environment > Agents** menu, then select the agent group to be updated.
2. In the **Version** section of the **Settings** tab, a message will inform you that a new version is available. Choose the new version to apply to agents in this group.
3. Click on **Save** at the top right of the window to save changes.
4. In **Security > Deployment**, click on **Deploy**.
   The new configuration will be applied to agents in the group the next time they connect to the agent handler.
   You can apply the update to the agent more quickly by clicking on **Check for updates** in the agent interface. For further information, refer to Understanding the agent interface on workstations.

7.7.3 Applying the update manually to an agent not connected to the agent handler

If your agent is not connected to the agent handler or if you wish to control your agent updates, you must generate an installer and run it manually on the agents, as you would do during initial deployment. For more information, see Installing agents on workstations.

When updating, not only is the new software version applied to an agent, but also the new configuration version, including security policies and agent group configuration.

To ensure a successful update:

- The updated agent must belong to the agent group for which the installer was generated,
- The version of the configuration (e.g., policies and agent group configuration) included in the update must be more recent than the version of the agent's configuration.

If you do not meet these conditions, force an update on the agent.

7.7.4 Forcing an update on agents

With a standard installer, the configuration of an agent group cannot be applied to agents that do not belong to this group. The installer also does not allow downgrades to an earlier configuration version. For this, you must perform a forced agent update. It is better for the agent to be disconnected from the agent handler when you force an update, because the next time the agent connects to the agent handler, the agent will go back to the group that was initially assigned to it.

1. Select the **Environment > Agents** menu, then select the agent group that you want to apply to the agent.
2. In the **Agents** tab, click **Installer > Forced update > Generate 64-bit installer**.
3. Save the *AgentSetup_x64.exe* installation file to the location of your choice and run it on the agent as in an initial deployment. For further information, refer to Installing agents on workstations.
4. If you want to prevent the agent from returning to its original agent group the next time it connects to the agent handler, move the agent to the desired group before it reconnects. For further information, refer to the section Moving agents from one group to another.

Force an update if agents in Maintenance mode need to be updated. For further information on Maintenance mode, refer to the section Understanding self-protection on agents and performing maintenance operations.

7.8 Managing a pool with agents in different versions

If you update SES Evolution to a new version, but some agents or agent groups stay in the older version, your pool will be disparate.

For further information on updating agents, refer to the section Updating agents.

You can easily view the number of agents for each version in the dashboard under the Agents diagram. This information can also be found in the Agents tab of the Agents panel.

In a disparate pool, agents that kept the older version will not be equipped with all the new features. Information about incompatibility is shown as icons and descriptive tool tips in the **Environment > Agents** menu of the administration console.

Icon	Meaning
[2.3+]	The software version selected in the agent group does not support this feature. This is the case, for example, with the **Allow agent uninstallation** parameter in the configuration of an agent group. However, you can save the configuration of the agent group and deploy the policy on the group.
[2.2+]	The software version selected in the agent group does not support any features. This is the case, for example, when applications are filtered by version 2.2 command line arguments. If the icon is red, it means that you cannot save the group's configuration or deploy the policy.
⚠	Some agents in the group are not yet equipped with the version required to use a feature. This is the case, for example, with the feature that automatically assigns agents to a group.

7.9 Removing obsolete agents from the console

When agents are no longer used on the company's workstations, they continue to appear in the monitoring table of the **Environment > Agents** panel in the console, and are counted in the number of agents that the license allows.

We recommend that you keep your agent list up to date to avoid exceeding the number of agents that your license allows, and populating the database with agents that no longer exist.

There are two ways in which you can clean up your list: automatically and periodically removing agents or merging duplicates.

The **Agent logs - Modify** privilege is required to delete obsolete logs.

7.9.1 Removing offline agents automatically

Offline agents can be deleted automatically. This feature is configured separately for each agent group and takes place at regular intervals. It addresses the following scenarios:

- When a workstation is remastered after an employee leaves the company, changes computers or because the workstation required an operating system update, for example.
- When a computer is no longer used in the company.
- When an agent was uninstalled on the workstation even though it was disconnected from the agent handler when it was uninstalled.

To schedule the periodic and automatic deletion of agents that have not connected to agent handlers for a specified duration:

1. Select an agent group in the **Environment > Agents** menu and click on **Edit** at the top right side.

2. In a group's **Status and logs** tab, go to the **Agent real time monitoring** section.

3. Set the number of days for the **Automatic deletion after** parameter. The default value is 30 days.

Automatic deletion tasks are launched at 2 a.m. The time cannot be changed.

If an agent that was deleted from the console attempts to connect again to its agent handler, a new identity will be assigned to it.

7.9.2 Merging duplicate agents

Duplicates are merged globally on all agents. This operation is manual and with instantaneous results. It addresses the following scenarios:

- When a workstation is remastered but the same computer name is kept.
- When you do not want to wait until the next automatic deletion of offline agents.

To merge duplicate agents:

1. Select **All agents** in the **Environment > Agents** menu.
2. Go to the **Maintenance** tab.
3. Select a **Criterion for duplicate display**:
   - **Active Directory name**: when all the workstations are in the Active Directory, the Active Directory name is the best criterion as it guarantees the uniqueness of agents, so any duplicates detected can be deleted.
   - **Computer name or NetBIOS name**: these criteria can be chosen if some of the workstations are not in the Active Directory, because in general these are unique names.
   - **IP address**: this criterion can be chosen when several hosts in the company's pool have the same names. However, several hosts may share the same IP address, so use this criterion with caution.
4. Select one or several lines. Each line shows both agents; the one that connected most recently is shown first.
5. Click on **Merge**. All grayed out agents will be removed from the database.

7.10 Uninstalling agents

There are several ways to uninstall an agent. Administrator privileges are required, and the operation must also be allowed in the agent group configuration, except for the challenge-based method. For further information, refer to the section Allowing administrators to uninstall agents.

- To uninstall several agents via a GPO, MECM (formerly SCCM), etc., run the executable EsSetupWorker.exe located in the installation_folder\SES Evolution\Agent\Bin. The default installation folder is C:\Program Files.
- To uninstall an agent from an individual workstation, use the **Uninstall** menu in **Programs and Features** in the Windows control panel. *EsSetup.exe* is the program that will be run. If you want to use it in a script, refer to the section on Using the EsSetup command to find out which options it offers.
- To uninstall an agent without administration privileges, using challenges, refer to the section Resolving issues with challenges.
- If none of the above methods work, you can force an uninstall with the *AgentRemovalTool.exe* tool. Refer to the section below on how to use the tool.

In any case, the uninstallation takes effect only after the workstation has been restarted. All files associated with the agent will be deleted except the registry key *HKEY_LOCAL_MACHINE\Software\Stormshield\SES Evolution* containing the agent's unique ID, which can be reused for a future installation.

The log file that captured the uninstallation will also be kept in the temporary folder of the user who uninstalled the agent.

7.10.1 Using the EsSetup or EsSetupWorker commands

The EsSetup command makes it possible to uninstall or repair the SES Evolution agent. The following are the options that this command offers:

`/silent or /s`	Hides the progress bar while the agent is being uninstalled or repaired.
`/Repair`	Launches an integrity check and an agent repair.
`/Log <path>`	Specifies the path of the agent's log file.
`/NewAgentId`	Deletes data regarding the agent's communication with the agent handler: unique ID, certificates used internally, as well as the ID and private data used in challenges. The agent retrieves new data the next time it connects to the agent handler.

Used without the `/Repair` or `/NewAgentId` options, this command uninstalls the SES Evolution agent.

The EsSetupWorker command runs the same options as EsSetup, but without a graphical interface. It is preferable for operations via GPO, SCCM, etc. It offers the same operations as EsSetup except for `-Silent`.

7.10.2 Forcing agent uninstallation

If standard uninstallation methods do not work, download the *AgentRemovalTool.exe* tool from your MyStormshield client area (32-bit and 64-bit versions available in the section **Downloads** > **StormshieldEndpoint Security** > **Evolution** > **Tools**). In command line, the tool allows you to force the agent to uninstall from the workstation in Windows safe mode.

To force the uninstallation of an agent, follow the steps below: Repeat the operation twice:

- The first time in Windows safe mode,
- The second time in normal Windows start mode.

Administration privileges are required to run the tool.

1. Start the workstation in safe mode.
2. If necessary, show help with the command `AgentRemovalTool.exe --help`.
3. Ensure that the tool detects the right version of the agent installed on the workstation by using the command `AgentRemovalTool.exe --supported-versions`.
4. Run the command `AgentRemovalTool.exe --remove` to start uninstalling the agent.

```
PS C:\tmp\AgentRemovalTool> .\AgentRemovalTool.exe --remove
Agent Removal Tool
This utility is a part of Stormshield Endpoint Security Evolution.
(C) Stormshield 2022

[INFO] Manual Uninstall script Menu

  1. Remove agent files, registry keys, and event logs
  X. Remove agent network objects (Windows Filtering Platform objects). NOT AVAILABLE: This cannot be run in safe mode
  3. Exit

Enter an option number:
```

5. Select menu number 1. Menu number 2 is not available in safe mode.
6. Next, start the workstation in "standard" mode.

7. Run the command `AgentRemovalTool.exe --remove` again.

```
PS C:\tmp\AgentRemovalTool> .\AgentRemovalTool.exe --remove
Agent Removal Tool
This utility is a part of Stormshield Endpoint Security Evolution.
(C) Stormshield 2022

[INFO] Manual Uninstall script Menu

  X. Remove agent files, registry keys, and event logs. NOT AVAILABLE: Windows must have been started in safe mode
  2. Remove agent network objects (Windows Filtering Platform objects)
  3. Exit

Enter an option number: _
```

8. Select menu number 2. Menu number 1 is not available in "standard" mode.
   Once you have completed these steps, the agent will be uninstalled correctly.

9. Select menu number 3 to quit the tool.

The following commands are also available in safe and standard modes:

- If the tool is unable to detect the version of the agent installed, the command `AgentRemovalTool.exe --remove --agent-version "2.3.2.0"` makes it possible to indicate which version to remove.

- If the agent failed to uninstall at the first attempt with the tool, the command `AgentRemovalTool.exe --remove --force` makes it possible to launch a second attempt.

7.11 Understanding the interactions between SES Evolution and Windows Defender

The SES Evolution agent automatically adds exceptions to Windows security system settings to prevent the Windows Defender antivirus program from triggering and blocking when SES Evolution performs legitimate operations.

These exceptions are added to the **Virus and Threat Protection** menu in the **Windows Security** window:

- in the list of controlled folder accesses,
- in the list of exclusions.

Thanks to these exceptions, Windows Defender does not block the following SES Evolution actions:

- Yara scans or Indicators of Compromise (IoC) searches,
- creation of Windows snapshots to protect against ransomware attacks. This creation is blocked if Windows Defender Ransomware protection is enabled.

These exceptions are removed when the agent is uninstalled.

> ℹ **NOTE**
> We recommend closing the **Windows Security** window when installing, updating or repairing a SES Evolution agent.

8. Managing security policies

A security policy is applied to SES Evolution agents to control access to resources and protect workstations against malicious behavior.

Before implementing security policies on your machines, you can test them transparently for users. You can then evaluate their impacts and adjust them if required. For more information, refer to the section Testing security policies.

8.1 Understanding security policies

A security policy consists of audit and protection rule sets. Each rule set is a set of security rules that applies to applications, ACL resources, network resources, devices and threat protection, which can be made private, i.e., specific to a policy, or shared among several policies.

Rule sets make it possible to pool rules for several policies, and manage various versions of these rule sets to create pre-production and production policies. Aggregating these rule sets in a single policy also makes it possible to load common rules over rules that are specific to your company's environment.

> ✏️ **EXAMPLE**
> You can alternate two policies based on a collaborator's location – one policy to manage access to internal resources, and one policy to manage access to resources when the collaborator logs in remotely. Both of these policies can share the same sets of rules, with only one differing set, so that they can block mobile devices from connecting to the network when they are not connected to their domain network. The different rule set allows these devices to log in to their domain via only VPN tunnels.

Once security policies are created, they will be linked to agent groups that will apply them to your pool. Only security policies can be linked to agent groups. Rule sets cannot be directly linked to agents.

You can test your policies before implementing them. For more information, refer to the section Testing security policies.

Security rules can be disabled at any time. For more information, refer to the section Disabling security rules.

8.1.1 Understanding built-in and custom security policies

SES Evolution allows the use of two types of security policies: built-in or custom.

Built-in security policies

SES Evolution is equipped with several built-in security policies that can block the behavior and techniques used by most malicious programs, regardless of their purpose, e.g., Trojan horses, remote control tools, ransomware, password stealers, etc. The following are built-in policies:

- **Simplified default policy** - enables the quick and simple deployment of SES Evolution in a pool by dedicating few human resources to it and without the need to modularly manage administration. Although it can be used without any specific configuration, you must still know how to operate the administration console to create exceptions and update policies.

- **Default policy** - constitutes a balanced compromise between the need for administration and the security level matching most organizations' needs. Targets companies with moderately large security teams, and which know how to handle SES Evolution administration databases.
  This security policy is applied by default to agent groups.
- **Hardened default policy** - raises the security level in a pool to the highest level, making administration harder. It is important that you test it with a pilot group before deploying the policy, to benefit from its policies while keeping false positives to a minimum. Used by companies with mature security teams and a well-defined security policy (e.g., an approved software catalog). It requires regular maintenance by administrators.
- **Backoffice component protection policy** - It guarantees the protection of SES Evolution backoffice components: the backend, agent handlers and the administration console. It contains protections from the default policy, but with the addition of several protection rules that strengthen the security of protected processes and block attempts to read or modify their configuration data.
  You can apply this policy as is to agent groups that contain backoffice components.

Built-in policies consist of built-in rule sets. For more information, refer to the section **Understanding built-in rule sets**.

Custom security policies

If built-in policies do not cover all use cases, you can create custom security policies that adapt closely to your infrastructure. To do so, use the rule sets that make up the built-in policies or create your own rule sets. For more information, refer to the section **Creating security policies**.

> ✎ **EXAMPLE**
> Create rules to manage access to the corporate network of your mobile collaborators, or manage the use of trusted devices in your pool.

8.1.2 Understanding the difference between protection, exception and audit rule sets

There are several types of rule sets: protection, exceptions, and audit.

They serve different purposes depending on the rule set to which the security rules belong.

- In a protection rule set, rules can be used to block attacks on workstations, detect elevation of privileges and manage access to different applications, networks, peripherals, etc.
- In an audit rule set, the rules can be used to generate logs for the sole purpose of monitoring the activity of your pool, and possibly for reconstructing the context of an attack.
- An exception rule set contains only exception rules. These are usually created from logs that you consider to be false positives. For further information, see **Adding exceptions for logs**.

The **Threats** tab of rule sets does not list exactly the same protections depending on whether it is a protection/exception set or an audit set. For more information, see the section **Managing vulnerability exploitation**.

Similarly, management of temporary web access and control of Wi-Fi board activation are only possible within the protection and exception rule sets.

Understanding protection and exception rule sets

In a protection or exception rule set, the agent evaluates the rules one by one and in the following order:

- If an action is prohibited for a resource, the agent will generate a log, block the action and stop scanning any other rules that apply to this resource.
- If an action is explicitly allowed for a resource, the agent will allow it and stop scanning any other rules that apply to this resource.
- If a rule does not apply to a resource, the agent will continue scanning the rules that follow.

Protection rule sets, are used to protect your workstations against malicious behavior and restrict access to protect your assets against dangerous behavior.

Exception rule sets allow you to adjust restrictions to limit false positives.

In protection and exception rule sets, all resource and device access control rules have a **Passive rule** mode. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode.

You can also test entire rule sets or entire policies before implementing them on your pool of machines. For more information, refer to the section Testing security policies.

Understanding audit rule sets

In audit rule sets, if **Audit** is selected as the action in a rule, the agent sends logs to indicate the actions performed by applications. The agent scans all the rules that follow in all cases.
Use this mode to monitor access to certain resources and send relevant information to the administrator without blocking access, so that abnormal behavior can be detected.

Audit rules can also be configured to monitor collaborators' activity: the applications that they use most often, or the versions of the applications that they use for example.

To prevent too many logs from being generated, create precise rules that do not cover too wide a range of resources or applications.

Audit rules can be used transparently in SES Evolution if you choose not to show logs on the agent or console, or if you choose not to send them to a syslog server. However, during an attack, the logs generated and saved on the agent can help to reconstruct the context of the attack, which is illustrated in a chart. For further information, refer to the section Analyzing contexts to understand attacks.

In audit rules, each action can be set to: **Allow** or **Audit**. **Allow** means that the rule will not do anything. It may be useful when you want to configure a default action and one or several specific actions in a rule. You can select **Audit** for specific actions and **Allow** for the default action. It is also useful when there are several actions available for a resource and you want to monitor only one type of action.

8.1.3 Organizing rules and rule sets in a policy

The agent evaluates protection and audit rules in the same order that rule sets appear in the policy, and in the same order as the rules inside these sets. If several rule sets apply to the same resources, ensure that the order of the rule sets is correct, as rules will no longer be evaluated once a rule is applied to the agent. This means that the rules highest up in the rule set are applied.

All rules in a policy, regardless of whether they belong to private or shared rule sets, are aggregated as if they were created in the same policy. If a policy contains two rule sets, all the rules from the first set will be read before the rules in the second set.

In general, if you use both protection and audit rule sets in the same policy, we recommend that you put audit rules before protection rules. This guarantees that logs will be generated for the actions that you want to monitor. If you put protection rule sets before audit rule sets, and both sets apply to the same resources, audit rules will not be read once a protection rule applies, and no audit logs will be generated.

Conversely, even when an audit rule applies, the agent continues to read rules, so protection rules will be evaluated.

In the event that you wish to create a policy comprising audit rule sets and protection rule sets provided by Stormshield in shared rule sets, and custom rule sets tailored to your environment, we advise you to take inspiration from the order of rule sets recommended in the Recommendations section of the *Release notes* SES Evolution.

For more information on rule sets provided by Stormshield, refer to the section Understanding built-in rule sets.

- Hover your mouse over a rule set to display the drag-and-drop icon on the left of the set and change the order.

The order of rules in the same protection rule set also matters, as they follow the same evaluation criteria as those in rule sets. Rules are evaluated in order and will stop being evaluated once a rule applies. Rules that apply to specific resources must therefore be placed before more general rules. The same goes for specific behavior in a rule. Refer to the next section for more information on specific behavior.

8.1.4 Using default behavior and specific behavior in rules

In access control rules, you can apply a default behavior or action and one or several specific actions.

When to add specific behavior

Define specific actions if you want to allow or block access from identified applications to the resource targeted by the rule.

You can add several specific actions in the same rule: one for example, to allow certain applications, and one to block others. In this case, the order of specific actions matters: if an application in the first specific action can access the resource, the rule applies and the second specific action will not be read.

When to enable default behavior

In a protection rule set, enable default behavior when you want to ensure that access to the resource will be blocked or allowed, regardless of which rules follow.

✏️ **EXAMPLE 1**
In code execution rules:
- To prevent the DLL *\lxssmanager.dll* from running on all applications, enable default behavior with **Execution = Block**.
- To prevent the DLL *\system.management.automation.dll* from running on all applications except legitimate applications, enable default behavior with **Execution = Block** and add a specific

action **Legitimate applications = Allow**.

> 📝 **EXAMPLE 2**
>
> In file access rules:
> To always allow the account *NT SERVICE\TrustedInstaller* to run powershell scripts (*.ps), enable default behavior with **Read = Allow**.

In the examples above, access will ALWAYS be allowed or blocked. Enabling default behavior means that rules for the resource in question will no longer be read. As such, any rules placed after this rule will not apply.

In audit rules, the default behavior is ignored, and all rules will be read every time.

When to disable default behavior

In a protection rule set, disable default behavior when you want to ensure that the next rule affecting the same resource will be read.

> 📝 **EXAMPLES**
>
> - In the file access rules, to ensure that different log levels are generated according to which applications access the same resource, disable default behavior on the first rule, add a specific action to block certain applications, and apply a specific log level.
>   Next, create a second rule with a different specific action and a different log level.
> - In the process access rules, create the first rule to grant the Windows task manager all permissions to all processes and disable default behavior. With this configuration, the task manager will never be blocked by subsequent rules that may prohibit some applications from accessing some processes that may include the manager.

8.2 Creating security policies

A security policy consists of audit and protection rule sets. Each rule set is a set of security rules, which can be made private, i.e., specific to a policy, or shared among several policies.

For further information on security policies, refer to Understanding security policies.

Several versions of policies or rule sets can coexist and you can choose which version to use at any time. For further information, refer to the section Managing versions of a policy or a rule set.

Before you create security rules for your policy, create application, driver and network IDs first. For more information, refer to the section Creating identifiers.

To build your policy based on the default built-in rules provided by Stormshield (i.e., Default Policy), or based on your own rules, follow the instructions below.

You must hold the **Policies-Modify** privilege to create and modify security policies and identifiers.

8.2.1 Understanding built-in rule sets

Stormshield provides a series of rule sets built into the console. Some are already contained in the built-in security policies. You can also use them in your own custom policies. For more information, refer to the section Creating a security policy.

To view built-in SES Evolution rule sets, select the **Security > Policies** menu and click the **Shared rule sets** link. Built-in rule sets are those that have the prefix Stormshield - .

Other rule sets that Stormshield provides are not built into the console and can be downloaded from your MyStormshield personal area, or from the Stormshield download server.

Rule sets are regularly updated. You will find the Release notes regarding these rule sets, which contain their descriptions, in your Mystormshield personal area and from the panel in which updates can be downloaded.

In addition, new rule sets are regularly added on the Stormshield update server and also published on the MyStormshield download area.

The sequence of rule sets in a policy matters. For more information, refer to Organizing rules and rule sets in a policy.

Built-in rule sets can neither be modified nor deleted.

8.2.2 Customizing built-in rule sets

Some built-in rule sets must be adapted to your environment before they can be used.

This is especially the case for the five II 901 rule sets. These rule sets are based on the ANSSI's French directive, drafted with the purpose of protecting sensitive information systems, These are templates that you can customize, to help you create security policies that apply the recommendations of the Interministerial instruction no. 901 in practice.

Since built-in rule sets are read-only, you must duplicate them in order to customize them, then add them to your policies.

To duplicate a built-in rule set:

1. Select the **Security > Policies** menu.
2. Click **Shared rule sets** at the top right side of the panel.
3. On the line of the rule set to duplicate, click on the ⋮ icon, then **Duplicate**. The duplicated rule set appears at the bottom of the list of rule sets with a number between brackets.
4. Double-click on the duplicated rule set, then click on **Edit**.
5. Rename the rule set then adapt the rules and IDs to your environment.
6. Add the rule set to your policies by following the procedure explained in the section Creating a security policy.

8.2.3 Creating shared rule sets

Shared rule sets make it possible to pool rules for several policies.

If you want to use shared rule sets in your security policies, you can create them earlier, either separately or directly in a policy.

If you are running in pre-production and production environments, you can test a private rule set in a pre-production policy and change it to a shared set once you are sure that it is effective, so that it can be used in a production policy.

To create a shared rule set separately from a policy:

1. Select the **Security > Policies** menu.
2. Click on **Shared rule sets** at the top right side of the panel.
3. Click on **Create**. The **Create a rule set** window appears.
4. Select the type of set and name it. For more information on rule set types, see **Understanding the difference between protection, exception and audit rule sets**.
5. Click on **Create**.
6. You are now about to create the rules for your rule set. Click on the new rule set and click on **Edit**.
7. Use the tabs **Threats**, **Application**, **ACL resources**, **Networks** and **Devices** to add security rules to your rule set. For further information on how to create rules, refer to the sections **Defining threat protection rules** and **Defining access control rules**.
8. Click on **Save** at the top right of the window to save changes.

To use the rule set in a policy, see **Creating a security policy**.

8.2.4 Creating a security policy

Audit rule sets and protection rule sets can be set up within the same security policy, for example when you are building pre-production and production policies.

You can create as many rule sets as you need. Rules from different categories can be created in the same set, or you can create a set for each rule category. The general panel of each policy shows how rule sets are built:
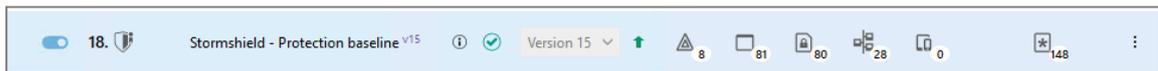
To create your own security policy:

1. Select the **Security > Policies** menu.
2. Click on **Create**. A line entitled *New policy* appears.
3. Double-click this line. The general panel of the new policy appears.
4. In the upper banner, click on **Edit**.
5. Enter a name and description for the policy. The description matters as it describes the various versions of the same policy.
6. In **Rule set**, click on **Add a shared rule set** to add an existing shared rule set, or on **Create a rule set** to add a new rule set.
7. If you are adding existing rule sets, select them in the order in which you want to see them appear in the policy. Their rank in the policy appears on the left next to their checkbox. For more information on the sequence of rules, refer to **Organizing rules and rule sets in a policy**.

8. If you are creating a new rule set, in **Create a rule set**:

   a. Select the type of set: Protection or Audit.

   b. Select who can see it: Private or Shared. Private sets are used only in the current policy. Shared sets can be used in several policies.

   c. Name the rule set.

   d. Click on **Create**.

9. You are now about to create the rules for your rule set. Click on the new rule set and click on **Edit**.

10. Enter a description of the rule set. The description matters as it describes the various versions of the same set.

11. Use the tabs **Threats**, **Application**, **ACL resources**, **Networks** and **Devices** to add security rules to your rule set. For further information on how to create rules, refer to the sections Managing vulnerability exploitation and Defining access control rules.

> 💡 **NOTE**
> Rules can also be copied and pasted between rule sets of the same type (audit or protection) and between policies.

12. In the general policy panel, you can change the order of rule sets by hovering over them with the mouse to display the drag-and-drop icon on the left. The order of the rule sets is important. For more information, refer to Organizing rules and rule sets in a policy.

13. Click on **Save** at the top right of the window to save changes.

For more information on versions of policies and rule sets, or if the 🔼 icon is displayed on the line), see Managing versions of a policy or a rule set.

Next, assign the security policy to the agent group you want this policy to apply to, then deploy it to your environment. For more information, see the sections Assigning a security policy to agents and Deploying the SES Evolution environment

8.2.5 Managing versions of a policy or a rule set

Several versions of policies or rule sets can coexist and you can choose which version to use at any time.

By managing several versions of a policy or rule set at the same time, you can set up pre-production and production policies and test how rule updates impact your pool. For example, your production policy can use a stable, i.e., tested and validated, version of rule sets while your pre-production policy uses a trial version that is more recent.

This feature also makes it possible to undo changes by redeploying an older version that worked correctly. E.g.: if you encounter a deployment issue in the environment, or if the deployment of a policy or rule set in your pool did not produce the expected results.

You can give your policies and rule sets accurate descriptions so that you can identify the various versions more easily.

When you export a policy or rule set, you export it in the version selected in the right side of the panel. For more information on importing and exporting policies and rule sets, refer to Importing and exporting policies and rule sets.

Managing various versions of a policy

In the general panel of a policy, version numbers are shown in the path of the policy at the top of the page, and in the right column. The last version deployed in your environment appears in blue. The version you are currently working on appears in green, or yellow if it is being edited.

After each policy deployment in your environment, editing the policy automatically increments its version number. You are therefore working on the latest version. The version of a deployed policy is always the latest version that was modified and saved.

For the latest version of the policy, successive changes are considered revisions of the same version of the policy. Click on a revision to go back to it at any time.

The [icon] icon indicates the revision you are currently working on.
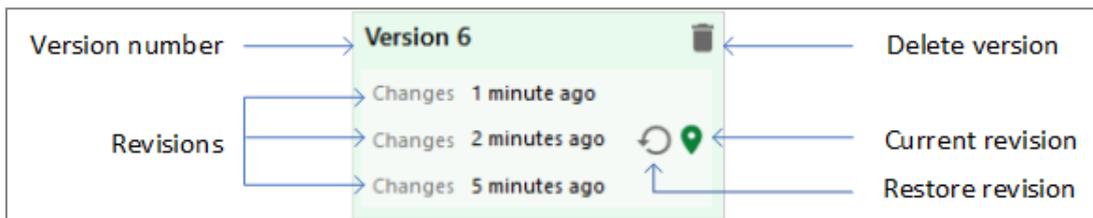
Only the latest version of a policy can be modified. Earlier versions must be restored before they can be modified.

Restoring a version of the policy:

1. Click on the desired version of the policy. The background will turn green.

2. Click on [icon] to restore this version. A new version will automatically be created with the content from this restored version, which therefore becomes the most recent. If the policy contains several revisions, you can restore a particular revision.

3. Make your changes and save them. If you deploy the policy in the environment, this is the version that will be deployed.

For more information on deploying a policy in your environment, see Deploying the SES Evolution environment.

Managing versions of a rule set

In the general panel of a rule set, version numbers are shown in the path of the rule set at the top of the page, and in the right column. The last version deployed in your environment appears in blue. The version you are currently working on appears in green, or yellow if it is in edit mode.

After each policy deployment in your environment, editing a rule set automatically increments its version number. You are therefore working on the latest version.

For the latest version of the rule set, successive changes are considered revisions of the same version of the policy. Click on a revision to go back to it at any time.

The [icon] icon indicates the version you are currently working on.

Only the latest version of a rule set can be modified. Earlier versions must be restored before they can be modified.

Restoring a version of a rule set:

1. Click on the desired version of the rule set. The background will turn green.

2. Click on ⟳ to restore this version. A new version will automatically be created with the content from this restored version, which therefore becomes the most recent. If the set contains several revisions, you can restore a particular revision.

3. Make your changes and save them.

Manually creating a new version of a rule set:

- Click on **Create new version** at the top on the right.

The **General** tab of a rule set shows the policies in which the rule set is used and the version number of the rule set for each policy.

Selecting the version of a rule set to use in a policy:

1. Go to the main panel of the policy:

2. Click on **Edit** in the upper banner.

3. Select the version for each rule set from the drop-down list.
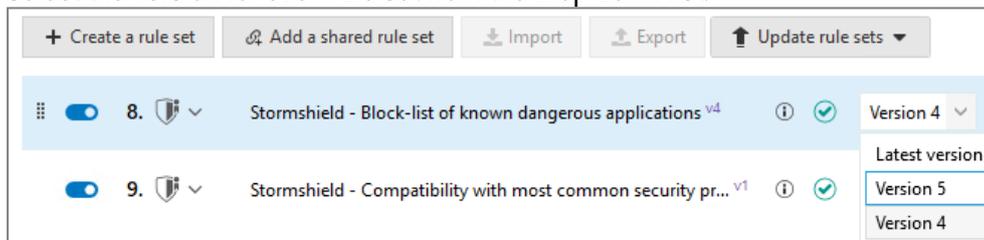
Multiple policies can therefore use various versions of the same rule set.

However, we recommend that you use a stable version of a rule set in your production environment.

If you have selected **Always use latest version** for a rule set, after a policy is deployed, the version number of the deployed rule set appears in the drop-down list. When you click on **Edit**, the **Always use latest version** parameter remains selected.

Updating policies with the latest version of a rule set:

Perform this operation only after the rule set has been tested and validated.

- In the **General** tab of a rule set, click ⬆ to update all policies using the same version of a set with the latest version of the rule set.

Deleting a version of a policy or a rule set

Versions of policies and rule sets can be deleted, including those provided by Stormshield.

However, a version currently being deployed, identified by the 🏢 icon, cannot be deleted.

1. Go to the main panel of a policy or rule set.

2. Click on the ✖ icon of the version that you wish to delete and confirm.
   When a version of the policy is deleted, all versions of private rule sets used in this version will also be deleted. No versions of shared rule sets will be deleted.

Automatically updating rule sets

The indicator ⬆ may appear in a policy's general panel:

- next to the version of a rule set [Version 1 ⌄ ⬆], when the selected version is not the most recent,

- next to the icon for rules regarding threats [⚠ 5], when at least one enabled advanced protection mode in the set is not using its latest version.

If you only wish to update the rule set in question:

1. Click on **Edit** in the upper banner.

2. If you wish to update each rule set individually, click on the action menu to the right of the rule set line, then on [⬆] to update advanced protection or the rule set as appropriate.
   - or -
   If you want to update all rule sets to the highest version, click the Update **rule sets** button. You can choose to perform the operation on all sets, or only on the Stormshield built-in sets.

8.3 Creating identifiers

Identifiers help to define the various applications, networks and drivers to which security rules apply. They are necessary when you create security rules, and must be created beforehand.

Each identifier consists of an unlimited number of entries linked by a logical "OR" operator, i.e., a security rule applies as soon as at least one of the identifier's entries is recognized.

Identifier entries make it possible to group various resources under the same identifier in order to pool any rules that may concern these resources.

There is no difference between creating two identifiers with single entry each and a single identifier containing two entries if all identifiers are associated with the same rule.

8.3.1 Creating application identifiers

Application identifiers, or application IDs, help to define which audit and protection rules apply to which applications, i.e.:

- Applications to protect or to exclude from a protection,

- Applications likely to interact with a protected application, for both legitimate or illegitimate purposes.

Since IDs are specific to each rule set, you must create IDs in each set. You can however export all the IDs of a rule set to import and use them in another set. For more information, refer to the section Importing and exporting identifiers.

> ✎ **EXAMPLE**
> If you want to prevent all applications from logging keystrokes on your web browser, except the virtualization tool, which has a legitimate need to log keystrokes. In this case, you need to create an application ID for the application you want to protect (web browser), and an ID for the legitimate keylogging application (virtualization tool).

Application IDs are necessary when you create rule sets, and must be created beforehand.

1. Select a policy in **Security > Policies**, then select a set of rules.

2. Click on the **Identifiers** tab at the top right, then on the **Application IDs** tab.

3. Click on **Edit** in the upper banner, then on **Add an ID**.
   A blank ID appears below the existing IDs.

4.   Click on **Edit** at the bottom right.

5.   In the field **New application ID**, enter an ID name, then a description if needed.

6.   Click on [+] and select all the identifier criteria you wish to use. E.g. **Path** and **Certificate**.

7. Click outside the criteria window and define each ID criterion selected:

Paths

a. Click on **Edit** then in the blue field at the bottom, enter the partial or full path to the application's executable file. This path may be a link or the path in the file system. The characters * and ? are allowed. Enter for example *\Apache.exe to identify the Apache application regardless of its location on the workstation.
Full paths beginning with a letter (i.e., E:\Data\Backup) are not supported if the **Volume type** is remote or removable.
Stormshield strongly recommends the use of EsaRoots path roots provided by SES Evolution instead of drive letters (i.e., C:\...). In fact, these letters may differ from one workstation to another.

b. You can also specify an alternate data stream. The ADSs of an executable file allow it to be looked up by several data streams. For further information, refer to Microsoft Windows documentation.

c. Click on **Add**.

d. Enter other paths in the blue field if necessary, then click on **Add**.

e. Click **OK** to confirm the path list.

Hashes

Hashes make it possible to accurately identify a trusted binary, as any modification will change the hash, which will no longer be recognized. Hash-based identification can be used in the following cases:

- To guarantee that a legitimate binary has not been replaced or modified. However, this requires tedious maintenance as you will need to change IDs after every software update. It should therefore be used only on systems that do not undergo many changes.

- To identify malware programs that often change names but may keep the same hash. Import the list of the most common malware hashes to block them from running.

To add hashes:

a. Click on **Edit** then on the pencil icon.

b. In the blue field at the bottom, enter the MD5, SHA1 or SHA256 hash of the application's binary file and a description, then click on **Add**.
To obtain the hash of a binary, you can use the following Powershell command. In this example, the SHA256 hash of all *.exe* files is obtained:

```
Get-ChildItem -Recurse -Filter '*.exe' | get-filehash -Algorithm
SHA256 | select path, Hash
```

c. Enter other hashes the blue field if necessary, then click on **Add**.

d. Click on **OK**.

e. You can also import a list of hashes from a CSV or text file. The file must contain one hash and a description per line separated by a comma, tab or semi-colon:

- Hash (MD5, SHA1 or SHA2),

- Description.

If there is an error or duplicate hash, SES Evolution will indicate it and only valid and unique hashes will be imported.
Once they have been entered and imported, the window shows the number of hashes for each algorithm.

f. Click on **OK** to confirm the list of hashes.

Parent process

Select the process that launches the application. The application ID of this process must be created beforehand.

a. Click on **Manage**.

b. Search for the ID of the parent process(es) using the search field and select them from the list that appears.

c. Click on **OK** to confirm the list of parent processes.

Certificate

Import the digital signature certificate provided by the application's vendor. Certificate-based identification can be used in the following cases:

- To strengthen the identification of a trusted binary by making it less restrictive than a hash because the certificate does not change with every new version of the binary. It is more reliable than using just a path because an attacker can always rename a malware program to *winword.exe* for example.

- To trust a vendor and therefore all software that it signs with its certificate. For example, you can allow the execution of all binaries signed by Microsoft, or even all binaries signed by a trusted certification authority.

To obtain a certificate, you can use the following Powershell command. In this example, we obtain the Acrobat Reader certificate, which we will name *Adobe.cer*:

```
(Get-AuthenticodeSignature -FilePath "C:\Program Files
(x86)\Adobe\Acrobat Reader
DC\Reader\AcroRd32.exe").SignerCertificate | Export-Certificate -
FilePath Adobe.cer
```

To add certificates:

a. Click on **Manage** to the right of the checkbox, then on **Import certificate**.

b. Choose the certificates to import.

c. Search for the certificates(s) using the search field and select them from the list that appears.

d. Click on **OK** to confirm the list of certificates.

e. The **Restrict search to the list ({0}** option makes it possible to search for the applications identified by the certificates listed in the ID. Keep the checkbox unselected (by default) if you want a rule to search among all certificates that have signed applications. In this case, certificates do not need to be imported in the identifier.
   You can refine the search by using the options **Valid** and **Not valid** so that a rule applies when it finds valid or invalid certificates (regardless of whether they are listed in the ID). Invalid certificates refer to certificates with the following status: not signed, revoked, untrusted, expired, corrupted or missing. These options are compatible with agents from version 2.3 upwards.

f. If needed, select **Ignore errors when validating the certificate signature** so that rules using this ID apply anyway when the certificate cannot be verified and an application is identified by other defined criteria in the ID.

The verification of the certificate may fail in the following cases:

- when a binary file starts running before the agent is installed, before the agent's services start, when the firewall restarts for the first time after the installation of the agent, or after a major update of the operating system,

- when the binary file *WerFault.exe* is involved, as it is processed differently from other files for technical reasons,

- when a third-party app prevents the binary file from being blocked.

Run in

a. Click on **Edit**.

b. From the drop-down list at the bottom, select the type of account that launches the identified application (e.g., *NT_AUTHORITY\System*), then click on **Add**.
You can choose a *strict* or *higher than or equal* integrity level. For example, if you set a Powershell identifier with execution context to *Administrator-level integrity (strict)*, the rule is triggered only for a Powershell run by an Administrator. On the other hand, with an execution context set to *Administrator-level integrity (or higher)*, the rule is triggered if Powershell is run by an Administrator or the Operating System.

To define a specific account, enter the account's SID (security identifier) in the field.

To obtain an SID, open a command window with administrator privileges and run the following command:

```
WMIC useraccount get name,sid
```

c. Select other accounts if necessary and click on **Add**.

d. Click on **OK** to confirm the list of accounts.

Volume type
Enable the volume type(s) on which this application runs: local disk on the workstation, network share (e.g., Samba/CIFS, DFS, etc.) or removable device (e.g., USB key, external hard disk, mobile phones depending on their configuration, etc.).

Command line
Filter applications based on arguments in their command line. making it possible to apply different rules to the same application, depending on how it is used. Refer to the next section for details on how to use the command line criterion.

Specifying more criteria will more accurately identify the application because all criteria must match.

> ✏️ **EXAMPLE**
> By specifying the application *PowerShell.exe* signed by *Microsoft*, launched by the scheduled task *schtasks.exe*, running from the local disk via the account *NT_AUTHORITY\System*, all five criteria must match for the application to be identified.

8. Click on **Add an entry** if you want to add another list of criteria for the same ID. Having several entries makes it possible to group various resources under the same ID, if the same security rules use them. For example, you can group various browsers together, or group various dangerous applications to set up a blacklist.

9. Enable the option **Include child applications of the applications identified below** so that when a rule is applied to an ID, it will also apply to all of its child applications. This option helps to identify installation programs that are extracted into a temporary folder and run executable files that have random names. By declaring the installation program a legitimate program, all the temporary files that it creates and launches will also be considered legitimate.

10. Click on **OK**.

11. If you have finished creating application identifiers, click on **Save** in the upper banner.

12. To show the contents of an application identifier without editing it, click on **View**.

> 💡 **TIP**
> Application identifiers can also be created directly from a rule. In a rule, click on

> $\boxed{+}$, then on **Create a new identifier**.
> Likewise, from a rule, you can click on a selected identifier to modify it. Changes will also apply
> to the identifier in the **Identifiers** tab.

Filtering applications via command line arguments

In application IDs, you can indicate command line arguments as an identifying criterion.

This criterion makes it possible to apply different rules to the same application, depending on
how it is used, to gain better control over the use of certain applications.

> ✎ **EXAMPLE**
> With this type of filtering, you can prevent PowerShell from running only when it is run as an
> invisible process, or when its command line arguments attempt to bypass Windows execution
> policies. Such behavior may indeed be deemed malicious.

Managing compatibility with agent versions

This feature can be used with agents in at least version 2.2.2. If a group of agents in a version
lower than 2.2.2 applies a policy that contains application IDs using the **Command line** criterion,
indicators will appear in various parts of the console to indicate incompatibility. For more
information, refer to the section Managing a pool with agents in different versions.

To make your agent pool support this feature, the SES Evolution release notes explain the
procedure of updating built-in security policies and agent pools in Recommendations.

Using the Command line criterion in an identifier

To create an application identifier based on command line arguments:

1. Select a policy in **Security > Policies**, then select a set of rules.
2. Click on the **Identifiers** tab at the top right, then on the **Application IDs** tab.
3. Click on **Edit** in the upper banner, then on **Add an ID**.
   A blank ID appears below the existing IDs.
4. Click on **Edit** at the bottom right.
5. In the field **New application ID**, enter an ID name, then a description if needed.
6. Click on $\boxed{+}$ and select **Command line**.
7. Click outside the criteria window.
8. Click on **Edit**.
9. Enter a name and select a mode:
   - **Custom parameters** (default mode): customize the parameters that the rule must search
     for in a command line.
   - **Contains at least one parameter**: the rule will apply every time it finds command lines
     containing at least one parameter.
   - **Without any parameter**: the rule will apply every time it finds command lines that do not
     contain any parameters.

10. If you have chosen **Custom parameters** mode, you will create one or several specifications by selecting options on the left, and by indicating command line parameters in the field on the right. If you are creating several specifications, they are linked by logical "AND" operators. This means that the rule using this identifier will apply if all the specified conditions are met.

a. Choose from the following options:

General

Not	The rule applies to all command lines that do not contain the parameter(s) indicated in the field on the right.
Case sensitive	The rule applies only when it finds the parameter(s) in the case specified indicated in the field on the right.
Command	After the specified parameter, the rest of the command line is interpreted as a nested command line. Such command lines are introduced with the `-Command` parameter for PowerShell or the `/c` parameter for cmd for example.

Parameter type

String	The parameter is a character string.
Flag	The parameter is found in a command line option beginning with `/` or `-` for example. To create an identifier corresponding for example to the editor of the registry that silently runs a *.reg* file, i.e. `regedit /s` in command line: 1. Create a **Path** criterion and enter `*\regedit.exe`. 2. Create a **Command line** criterion, select the **Flag** and **Is equal to** options, then enter `s` in the field on the right. The `/` or `-` characters do not need to be entered. Take note that the double dash `--` is not supported. For example, to search for the `--arg` argument, you must select **String** as the type of parameter.

Check

Is equal to	The parameter must be the same as the character string indicated in the field on the right.
Begins with	The parameter must begin with the character string indicated in the field on the right.
Ends with	The parameter must end with the character string indicated in the field on the right.
Contains	The parameter must contain the character string indicated in the field on the right.
Is prefixed with	The value that the rule recognizes rule may be a prefix of the character string indicated in the field on the right. For example, the character string "version" will match the values "v", "ve", "ver", etc., up to "version".

Position (visible from the second specification onwards)

None	There is no position criterion.
Followed by	The parameter searched for follows the previous parameter.
Immediately followed by	The parameter searched for immediately follows the previous parameter.

b. Enter one or several parameters in the field on the right. In the same specification, parameters are linked by logical "OR" operators. This means that the rule using this identifier will apply if at least one of the specified conditions is met.

11. When you have created all the specifications, confirm the creation of the "Command line" criterion.

12. Confirm the creation of the identifier.

Use case

As part of SES Evolution's anti-ransomware protection mode, this criterion type makes it possible to set **Process creation** rules on applications that may attempt to delete Windows shadow copies, among other operations. However, these shadow copies must be protected so that files encrypted by a ransomware program can be retrieved. For further information, refer to Managing ransomware attacks. These rules are included in the built-in Anti-ransomware protection rule set.

> ✎ **EXAMPLE**
> The use of VSSAdmin to manage Windows shadow copies can be allowed in your pool, except when it attempts to delete a shadow copy, for example. Indeed, a ransomware program may carry out such an action.
> In this case, create an application ID by indicating the following values for the **Path** criterion:
>
> | | |
> |---|---|
> | \EsaRoots\SystemRoot\system32\vssadmin.exe | 🗑 |
> | \EsaRoots\SystemRoot\syswow64\vssadmin.exe | 🗑 |
> | \EsaRoots\SystemRoot\WinSxs*\vssadmin.exe | 🗑 |
>
> Enter a path   Enter an alternate data stream   + Add
>
> ✓ OK   ✗ CANCEL
>
> Next, indicate the following values for the Command line criterion:
>
> Name: vssadmin - Shadow copies deletion   Mode: Customized
>
> + Add
>
> ↑ ∧ If string equals to delete   🗑
>
> **General**   **Parameter type**   **Check**   delete ✗
> ☐ Exclude   ⦿ String   ⦿ Is equal to   Add an argument
> ☐ Case sensitive   ○ Flag   ○ Begins with
> ☐ Command ❶   ○ Ends with
> ○ Contains
> ○ Is prefixed with ❶
>
> And
>
> ↑ ∧ If immediately followed by string equals to shadows   🗑
>
> **General**   **Parameter type**   **Check**   **Positioning**   shadows ✗
> ☐ Exclude   ⦿ String   ⦿ Is equal to   ○ None   Add an argument
> ☐ Case sensitive   ○ Flag   ○ Begins with   ○ Followed by
> ☐ Command ❶   ○ Ends with   ⦿ Immediately followed by
> ○ Contains
> ○ Is prefixed with ❶
>
> The ID will then include the following entry:
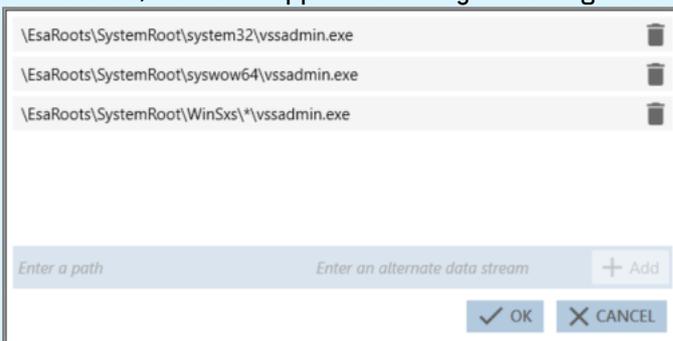
The ID can then be used in a blocking Process creation rule that kills the VSSAdmin application when it detects an attempt to delete a Windows shadow copy.

8.3.2 Creating driver identifiers

Driver identifiers, or IDs, make it possible to define legitimate drivers that you can exclude from rootkit detection.

Driver IDs are necessary when you create audit rules for rootkit detection, and must be created beforehand.

For more information, refer to the section Protection against various threats.

Since IDs are specific to each rule set, you must create IDs in each set. You can however export all the IDs of a rule set to import and use them in another set. For more information, refer to the section Importing and exporting identifiers.Importing and exporting identifiers

1. Select a policy in **Security > Policies**, then select a set of rules.
2. Click on the **Identifiers** tab at the top right, then on the **Driver IDs** tab.
3. Click on **Edit** in the upper banner, then on **Add an ID**.
   A blank ID appears below the existing IDs.
4. Click on **Edit** at the bottom right side of the entry.
5. In the field **New driver ID**, enter an ID name, then a description if needed.
6. Click on ⊞ and select all the ID criteria that you wish to use, e.g., **Path** and **Hashes**.

7. Click outside the criteria window and define each ID criterion selected:

Paths

a. Click on **Edit** then in the blue field at the bottom, enter the partial or full path to the driver file. This path may be a link or the path in the file system.
The characters * and ? are allowed. For example, enter *\drivers\Stormshield Endpoint Security Agent\es*.sys to include Stormshield drivers.
Full paths beginning with a letter (i.e., E:\Data\Backup) are not supported if the **Volume type** is remote or removable.
Stormshield highly recommends using the EsaRoots path roots provided in SES Evolution instead of drive letters (i.e., C:\...), as these letters may vary from one workstation to another.

b. You can also specify an alternate data stream. A file's ADS contains metadata and makes it possible to find out the origin of the file. For further information, refer to Microsoft Windows documentation.

c. Click on **Add**.

d. Enter other paths in the blue field if necessary, then click on **Add**.

e. Click on **OK** to confirm the list of paths.

Hashes

Hashes make it possible to accurately identify a trusted driver, as any modification will change the hash, which will no longer be recognized. Hash-based identification can be used in the following cases:

- To guarantee that a legitimate driver has not been replaced or modified. However, this requires tedious maintenance as you will need to change IDs after every software update. It should therefore be used only on systems that do not undergo many changes.

- To identify malware programs that often change names but may keep the same hash. Import the list of the most common malware hashes to block them from running.

To add hashes:

a. Click on **Edit** then on the pencil icon.

b. In the blue field at the bottom, enter the MD5, SHA1 or SHA256 hash of the driver and a description, then click on **Add**..
To obtain the hash of a binary, you can use the following Powershell command. In this example, the SHA256 hash of all .sys files is obtained:

```
Get-ChildItem –Recurse -Filter '*.sys' | get-filehash –
Algorithm SHA256 | select path, Hash
```

c. Enter other hashes the blue field if necessary, then click on **Add**.

d. Click on **OK**.

e. You can also import a list of hashes from a CSV or text file. The file must contain one hash and a description per line separated by a comma, tab or semi-colon:

- Hash (MD5, SHA1 or SHA2),

- Description

If there is an error or duplicate hash, SES Evolution will indicate it and only valid and unique hashes will be imported.

Once they have been entered and imported, the window shows the number of hashes for each algorithm.

f. Click on **OK** to confirm the list of hashes.

Owner

    a. Click on **Manage**.

    b. From the drop-down list at the bottom, select the type of account that launches the identified driver (e.g., *NT_AUTHORITY\System*), then click on **Add**.
To obtain an SID, launch a command window with administration privileges and run the following command:
```
WMIC useraccount get name,sid
```

    c. Select other accounts if necessary and click on **Add**.

    d. Click on **OK** to confirm the list of accounts.

    Specifying more criteria will more accurately identify the driver because all criteria must match.

8. Click on **Add an entry** if you want to add another list of criteria for the same ID. Having several entries makes it possible to group various resources under the same ID, if the same security rules use them. For example, you can group all legitimate drivers to compile a whitelist.

9. Click on **OK**.

10. If you have finished creating driver identifiers, click on **Save** in the upper banner.

11. To show the contents of a driver ID without editing it, click on **View**.

8.3.3 Creating network identifiers

Network IDs make it possible to define the network resources that you want to protect: IP addresses, ports, IP address ranges, or port ranges.

Network IDs are necessary when you create network rules, and must be created beforehand.

Since IDs are specific to each rule set, you must create IDs in each set. You can however export all the IDs of a rule set to import and use them in another set. For more information, refer to the section Importing and exporting identifiers.Importing and exporting identifiers

For more information, refer to the section Controlling network access.

1. Select a policy in **Security > Policies**, then select a set of rules.

2. Click on the **Identifiers** tab at the top right, then on the **Network IDs** tab.

3. Click on **Add an ID**.
A blank ID appears.

4. Click on **Edit** at the bottom right side of the entry.

5. In the field **New network ID**, enter an ID name, then a description if needed.

6. If you want the network ID to include all IP addresses EXCEPT the ones specified, enable the option **Invert identifier scope**.

7. The ID includes all IPv4 and IPv6 addresses by default. To specify certain addresses in particular, click on **No addresses added** and manually enter the values in the text field that appears. You can also add a description if necessary.

- To add several addresses at one go, separate them with commas in the text field and press Enter. Example: 192.168.128.254,192.168.95.15.

- To add an address range, separate the first value and last value with a dash and press Enter. Example: 192.168.131.0-192.168.131.100.

8. Click on **Finish changes**.

9. If you have finished creating application identifiers, click on **Save** in the upper banner.

8.3.4 Using path roots in identifiers

The workstations in your SES Evolution environment do not all have the same Windows installation. For example, the user profile and applications may be located in different drives from one workstation to another. SES Evolution provides variables in the form of path roots that allow rules to be adapted to each user, regardless of their drive names and trees.

Stormshield highly recommends the use of such roots in the **Path** field during the creation of application identifiers and file rules, especially to identify applications found in the *Programs* or *System32* folder.

Use the root...	To reference...
\EsaRoots\SystemDrive	The volume on which Windows is installed, typically C:
\EsaRoots\SystemRoot	The Windows folder, typically C:\Windows
\EsaRoots\UserProfiles	The Users folder
\EsaRoots\ProgramData	The folder in which applications automatically store data regardless of the user
\EsaRoots\ProgramFiles \EsaRoots\ProgramFilesX86	Folders in which 64-bit and 32-bit applications are installed respectively. On a 32-bit operating system, both symbolic links point to the same location.

> **EXAMPLE 1**
> Use the paths \EsaRoots\ProgramFiles\Internet Explorer\iexplore.exe and\EsaRoots\ProgramFilesX86\Internet Explorer\iexplore.exe to **create the application identifier** of the Microsoft Internet Explorer browser.
>
> Microsoft Internet Explorer [native]
>
> Paths : 1   ✕
>
> \EsaRoots\ProgramFiles\Internet Explorer\iexplore.exe
>
> View more
>
> Microsoft Internet Explorer [WOW64]
>
> Paths : 1   ✕
>
> \EsaRoots\ProgramFilesX86\Internet Explorer\iexplore.exe
>
> View more
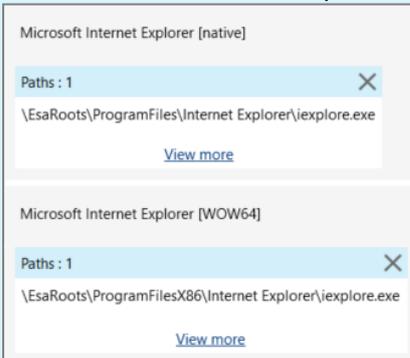
> **EXAMPLE 2**
> Use the path \EsaRoots\SystemRoot\System32\drivers\etc\hosts to identify the *hosts* file when

creating a file access rule.

8.3.5 Importing and exporting identifiers

The identifiers of applications, drivers and networks can be exported to a *.json* file that can be re-imported later. This makes it possible to:

- Use identifiers created for a rule set in a different rule set without the need to create them again.
- Transfer the list of identifiers to SES Evolution's technical support to make it easier to debug issues.

You can import/export lists of application, driver and network identifiers separately. However, you cannot select only some identifiers from the same list. They will all be imported/exported.

Exporting a list of identifiers

1. Select a policy in **Security > Policies**, then select a set of rules.
2. Click on the **Identifiers** tab at the top right, then on the **Application IDs, Driver IDs or Network IDs** tab.
   The list of IDs appears.
3. Click on **Export IDs** and choose the name of the *.json* file and the folder to which you want to export the file. All the IDs on the list will be exported.

Importing a list of identifiers

1. Select a policy in **Security > Policies**, then select a set of rules.
2. Click on the **Identifiers** tab at the top right, then on the **Application IDs, Driver IDs or Network IDs** tab.
   The list of IDs appears.
3. Click on **Import IDs** and choose the **.json** file you want to import.

8.4 Managing vulnerability exploitation

Hackers use many malicious techniques such as heap spraying and process hollowing to exploit vulnerabilities on workstations. Threat protection rules on Stormshield Endpoint Security Evolution make it possible to detect these attack techniques and/or block them effectively.

Depending on the severity of threats, some protections are available only in audit rule sets or only in protection rule sets, while some are relevant in both cases.

In protection rule sets, contexts are always generated for most rules. In audit rule sets, this is an option that you can choose to enable or disable.

The Stormshield Default Policy implements a specific number of protection and audit rules, but you can create your own custom rules. For every rule type, you can define:

- Default behavior,
- Specific behavior for certain applications.

For more information on audit and protection rule sets, and default and specific behavior, refer to Understanding security policies.

Security rules can be disabled at any time. For more information, refer to the section Disabling security rules.

8.4.1 Protection against various threats

SES Evolution provides rules that help you to detect the main threats and protect yourself from them. This section briefly describes the particularities of each type of threat and the corresponding protection measures. Refer to Configuring threat protection for information on how to implement protection against the various threats.

Built-in protections

SES Evolution provides rules that help you to detect the main threats and protect yourself from them. In this section, we briefly explain the characteristics of each type of threat. Refer to Configuring threat protection for information on how to implement protection against the various threats.

Process hollowing

The process hollowing protection mechanism detects and blocks malicious executables that attempt to disguise themselves as legitimate processes on the system (e.g., explorer.exe) so that they can run without being detected by Windows. It counters attacks such as RunPE and Process Doppelgänging.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always
Recommendations	Enable this protection by default in **Detect only** mode, and only disable it for well-identified internal applications that legitimately use the process hollowing technique.

Stack pivoting

Stack pivoting attacks exploit buffer overflows so that they can hijack an application's execution flow to make a legitimate application run malicious code.

The stack pivoting protection mechanism regularly monitors memory. If SES Evolution detects abnormal behavior on an agent, especially a different stack address, it will stop the process to prevent the code from being executed.

Rule set type	Protection
Log level	Alert
Generate a context	Always

Recommendations	Enable this protection by default in **Detect only** mode for all applications.

Execution flow hijacking

The execution flow hijacking protection mechanism detects and neutralizes malicious shellcodes that exploit buffer overflows to use the addresses of system functions in the dynamic library kernel32.dll.

Rule set type	Protection
Log level	Error
Generate a context	Always
Recommendations	Enable this protection by default in **Block and interrupt** mode for all applications. To avoid false positives, fine-tune the protection by declaring callers (.exe, .dll) and/or functions that can be explicitly authorized or blocked by the rule. For further information, see Case of protection against execution flow hijacking.

Heap spray

Heap spraying is a technique that consists of allocating large amounts of memory to facilitate the execution of malicious code after a vulnerability is exploited. Since heap spraying can only be used on 32-bit applications, the SES Evolution protection mechanism is not enabled on 64-bit applications.

Rule set type	Protection
Log level	Alert
Generate a context	Always
Recommendations	Enable this protection by default in **Block and interrupt** mode for all applications.

Access token manipulation

The operating system assigns a security token to every process; among other data, this token contains the account with which the process was run and the privileges associated with this process.

Some attack techniques manage to steal or duplicate the security tokens of high-privilege processes, thereby gaining access to resources or privileges that would not normally be granted to them.

The token protection mechanism on SES Evolution makes it possible to block such attacks by stopping the process that stole the token.

Rule set type	Protection
Log level	Alert
Generate a context	Always
Recommendations	Enable this protection by default in **Block and interrupt** mode for all processes.

Application hooking (Windows Hooks)

The Windows SetWindowsHookEx API allows a program to be notified when certain events occur on the system or on applications, e.g., mouse movements, keystrokes, etc. A DLL is injected

into target applications for this purpose.

Even though this is a legitimate mechanism, hackers may use it to inject malicious code so that a user's operations can be observed, e.g., the keystrokes when the user enters various passwords.

Rule set type	Protection and Audit
Log level	Protection: Error Audit: Information
Generate a context	Up to user (Yes by default)
Recommendations	Disable default protection.

When this rule is enabled, it controls all applications that use SetWindowsHookEx. If you do not want to completely block access to this API, do not enable this rule, but make the necessary adjustments in the Keylogging application rule.

Privilege escalation

This protection mode makes it possible to monitor applications' attempts to escalate privileges by using the Debug privilege. When this mode is enabled, SES Evolution compares the privileges usually granted to the application with those requested. If the requested privileges are higher, SES Evolution will consider the request a privilege escalation and may block the action.

Rule set type	Protection and Audit
Log level	Protection: Error Audit: Information
Generate a context	Up to user (Yes by default)
Recommendations	Up to user

EDR detection bypass

This protection mode protects against attacks from malicious programs that attempt to disable EDR (Endpoint Detection and Response) modules based on AMSI and ETW technologies.

Rule set type	Protection
Log level	Alert
Generate a context	Always
Recommendations	Enable this protection by default in **Detect only** mode for all processes.

Fileless attack

Fileless attacks act without writing malicious files to workstation disks. The attack occurs in memory.

This protection mode protects against processes that attempt such attacks.

Rule set type	Protection
Log level	Alert
Generate a context	Always
Recommendations	Enable this protection by default in **Detect only** mode for all processes.

Rootkit detection

A rootkit is a program that modifies the behavior of the operating system so that the system does not notice this program has been executed. Its aim is to gain and keep access to a computer, usually with malicious intentions.

Rootkit detection on SES Evolution makes it possible to monitor driver loading and verify their integrity.

Rule set type	Audit
Log level	Emergency
Generate a context	Up to user (Yes by default)
Recommendations	Enable these rules by default and disable them only for legitimate drivers.

Driver loading

The driver loading protection mechanism detects drivers that the operating system loads and generates a log for each driver.

Driver integrity

The driver integrity protection mechanism regularly verifies every driver to ensure that its integrity has not been potentially compromised, i.e., whether its major function table has been modified. If changes are detected, SES Evolution will identify the driver behind the attack and generates a log. For example, if a malicious driver could modify an antivirus driver, it would prevent files from being analyzed.

However, some drivers make legitimate changes, as is the case with some virtualization tools. These drivers must be excluded from the audit rule.

Advanced protections

Stormshield also provides a set of advanced protections against some types of threats. These protections are natively built into the administration console.

Advanced protections make it possible to detect and block malicious behavior on SES Evolution agents. They are based on heuristic analyses, which can be updated without the need to update the SES Evolution software.

To view advanced protections in the console:

1. Select the **Security > Policies** menu.
2. Click on **View advanced protections** at the top right side of the home panel of the policies.

Refer to Configuring threat protection for information on how to implement advanced protection against the various threats.

Advanced protections have version numbers and can be updated via Stormshield when necessary. During updates, you can therefore re-import them in the **Advanced protections** panel. All previous versions of a protection remain available in the administration console.

Kerberos ticket protection

Prevents the retrieval of Kerberos tickets from memory, as they may be used later to launch pass-the-ticket attacks.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Protection against ARP spoofing

Prevents network traffic from being intercepted, modified or stopped through ARP spoofing attacks. The ARP table is evaluated every 5 minutes.

Rule set type	Audit
Log level	Alert by default
Generate a context	Up to user (Yes by default)

WMI Persistence

This protection prevents malware programs from persisting on computers through WMI (Windows Management Instrumentation).
It relies on the *Microsoft-Windows-WMI-Activity/Operational* event log. In Windows 7 and Server 2008, the Windows update KB3191566 is needed for this log to be present.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Protection against malicious use of certutil

This protection mode protects users from the malicious use of the Windows program certutil, which allows certificates to be managed. Using this protection may generate false positives, as the files that certutil handles need to be opened in read-only mode. If such files cannot be accessed due to insufficient privileges, the operation on the certificates will be considered malicious, even though it is legitimate.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Environment discovery

This protection prevents the use of the built-in Windows tools that collect information on the host and system with the aim of performing malicious operations.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Ransomware

This protection mode keeps track of when files are modified and encrypted. If a particular number of such events occurs in the space of three seconds, the process in question will be stopped. This mode also makes it easier to retrieve data that the ransomware encrypts, by enabling:

- the identification of files modified by the ransomware,
- the restoration of the identified files, based on Windows shadow copies.

Rule set type	Protection

Log level	Alert by default
Generate a context	Always

Parent PID Spoofing

This protection mode prevents hackers from starting programs that they would declare as children of arbitrarily chosen existing processes with the purpose of concealing malicious processes from security analysts.

Rule set type	Protection
Log level	Critical by default
Generate a context	Up to user (Yes by default)

Advanced protections

Stormshield also provides a set of advanced protections against some types of threats. These protections are natively built into the administration console.

Advanced protections make it possible to detect and block malicious behavior on SES Evolution agents. They are based on heuristic analyses, which can be updated without the need to update the SES Evolution software.

To view advanced protections in the console:

1. Select the **Security > Policies** menu.
2. Click on **View advanced protections** at the top right side of the home panel of the policies.

Refer to Configuring threat protection for information on how to implement advanced protection against the various threats.

Advanced protections have version numbers and can be updated via Stormshield when necessary. During updates, you can therefore re-import them in the **Advanced protections** panel. All previous versions of a protection remain available in the administration console.

Sigma advanced protection

The Sigma format is a standard unified language for describing log-based incident detection rules. You can import Sigma rules into SES Evolution via the API or via scripts. For further information, see Importing Sigma security rules.

Rule set type	Passive protection (no blocking)
Log level	Depends on imported rule
Generate a context	No

ARP Spoofing

Prevents network traffic from being intercepted, modified or stopped through ARP spoofing attacks. The ARP table is evaluated every 5 minutes.

Rule set type	Audit
Log level	Alert by default
Generate a context	Up to user (Yes by default)

Parent PID Spoofing

This protection mode prevents hackers from starting programs that they would declare as children of arbitrarily chosen existing processes with the purpose of concealing malicious processes from security analysts.

Rule set type	Protection
Log level	Critical by default
Generate a context	Up to user (Yes by default)

WMI Persistence

This protection prevents malware programs from persisting on computers through WMI (Windows Management Instrumentation).
It relies on the *Microsoft-Windows-WMI-Activity/Operational* event log. In Windows 7 and Server 2008, the Windows update KB3191566 is needed for this log to be present.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Kerberos ticket

Prevents the retrieval of Kerberos tickets from memory, as they may be used later to launch pass-the-ticket attacks.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Environment discovery

This protection prevents the use of the built-in Windows tools that collect information on the host and system with the aim of performing malicious operations.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

Protection against malicious use of certutil

This protection mode protects users from the malicious use of the Windows program certutil, which allows certificates to be managed. Using this protection may generate false positives, as the files that certutil handles need to be opened in read-only mode. If such files cannot be accessed due to insufficient privileges, the operation on the certificates will be considered malicious, even though it is legitimate.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

<u>Ransomware</u>

This protection mode keeps track of when files are modified and encrypted. If a particular number of such events occurs in the space of three seconds, the process in question will be stopped. This mode also makes it easier to retrieve data that the ransomware encrypts, by enabling:

- the identification of files modified by the ransomware,
- the restoration of the identified files, based on Windows shadow copies.

Rule set type	Protection
Log level	Alert by default
Generate a context	Always

8.4.2 Configuring threat protection

The security rules that Stormshield provides include audit or protection rules that you can configure to protect your network from major attack classes that threaten workstations.

For more information on the attacks that SES Evolution thwarts, refer to the section Protection against various threats.

All threat protection rules are disabled by default. If there are several protection rule sets in your security policy, ensure that you enable the policy only for the set(s) in which you want to configure threat protection, and arrange your rule sets in the right order in the policy. If you configure threat protection in a rule set near the top of the policy, this rule may overload and cancel the effect of the threat protection configuration in the rule sets that follow.

Requirements

- For Driver loading and Driver integrity audit rules, you must create a driver ID beforehand for every legitimate driver to ignore.
For more information, refer to the section Creating driver identifiers.
- For all other protection types, application identifiers must be created beforehand for every application to be protected and for every approved application to be excluded from the protection rules.
For more information, refer to the section Creating application identifiers.

Configuring built-in protections

1. Select the **Security > Policies** menu and double-click on your policy.
2. Select the rule set in which you wish to add your rule.
The main page of the rule set appears.
3. Click the **Threats > Built-in Protections** tab.
Each protection type is a group of rules and can contain one or more rules.
4. If you are in read-only mode, click on **Edit** in the upper banner.
5. Enable the desired rule by clicking on the [ ] button on the left and click on the arrow to display its contents.

6. In the **Status** field in **Default behavior,** there are three or four statuses for each protection mode. Select:

   - **Allow**: SES Evolution does not block malicious actions and does not generate any logs.
   - **Detect only**: As in audit mode, SES Evolution detects malicious actions without blocking them, and generates logs for the administrator. But unlike audit mode, this option stops evaluating the rules that follow, and ignores them.
   - **Block**: SES Evolution blocks malicious actions and generates logs for the administrator.
   - **Block and kill**: SES Evolution blocks malicious actions and shuts down the process that launched the action.
   - **Block, kill and quarantine** SES Evolution blocks malicious actions, shuts down the process that launched the action and quarantines suspicious files. See Managing file quarantine.
   For audit rules, the available actions are always **Allow**, which does not do anything, and **Audit**, which generates a log and evaluates the next rule.

7. Click on **+ Add a specific behavior** to add the IDs of the applications for which the protection must behave differently. In *process hollowing* for example, you can enable the protection by default, and disable it specifically for your internal applications, such as virtualization tools, that use this operating mode.

8. In the upper banner in the rule, you can:

-   If necessary, rearrange the order of the rules by clicking on ⬚ when the cursor hovers the rule. Each rule displays its line number in the banner.

-   Disable rule. For more information, refer to the section Disabling security rules.

-   Indicate the intent of the rule, according to predefined categories:

-   ○ Unclassified: unclassified rule.

-   ○ Nominal: non-blocking rule conforming to nominal application behavior.

-   ○ Protect: blocking rule with a high log severity level.

-   ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

-   ○ Detect: non-blocking audit rule or passive rule.

-   ○ Context: rule used to build an attack graph.

-   ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

-   ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

    Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

-   Enter a description to explain what this rule aims to achieve.

-   Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
    Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

-   Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

-   Adding a comment.

-   Select the log settings that this rule will send.

-   Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

-   Deleting the rule.

9. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

10. Click on **Save** at the top right of the window to save changes.

Case of protection against execution flow hijacking

To avoid false positives due to this protection, you can fine-tune it by declaring executables or DLLs (callers) and/or functions that can be authorized or blocked.

> ✎ **EXAMPLE**
> Some streaming platforms use DLLs and features to control user access. These can be considered illegitimate by the SES Evolution execution flow hijacking protection. Declare them in the rule to authorize them.

1. Follow the Configuring built-in protections procedure.

2. To the left of the rule, click **Callers +,** and enter the path to one or more executable resources or DLLs. The generic characters "?" and "*" are allowed in this field.

3. Click on **Functions +** to select one or more functions.

To implement the example described above, you must:

- Allow the *Chrome* process to access the *GetFinalPathNameByHandleW* function via the *winevinecmd.dll* DLL,
- Block all other types of attempts (default behavior).
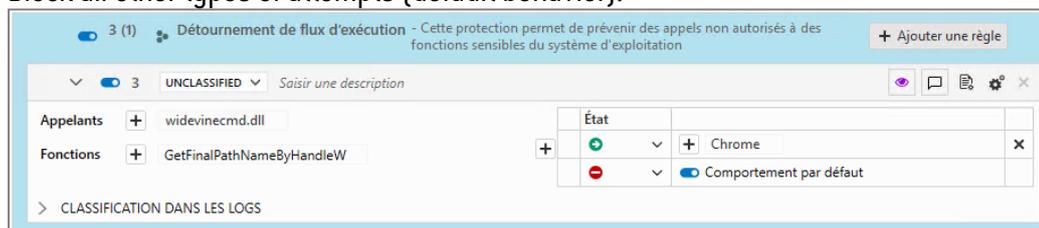
Configuring advanced protection

For more information, refer to the section Protection against various threats.

To enable and configure advanced protection:

1. Select the **Security > Policies** menu and double-click on your policy.

2. Select the rule set in which you wish to add your rule.
   The main page of the rule set appears.

3. Click the **Threats > Advanced Protections** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Enable the desired rule by clicking on the ⬜ button on the left and click on the arrow to display its contents.

6. In the **Version** drop-down list, indicate the version of the protection that you wish to run - either a version in particular or **Always use latest version**. If you are not using the latest version available, the indicator ⬆ will appear to the right of the drop-down list. It also appears on a policy's general panel when at least one enabled advanced protection mode is not using its latest version.

7. In the **Status** field, several statuses are available for each protection. Select:

   - **Allow**: SES Evolution does not block malicious actions and does not generate any logs.

   - **Detect only**: As in audit mode, SES Evolution detects malicious actions without blocking them, and generates logs for the administrator. But unlike audit mode, this option stops evaluating the rules that follow, and ignores them.

   - **Block**: SES Evolution blocks malicious actions and generates logs for the administrator.

   - **Block and kill**: SES Evolution blocks malicious actions and shuts down the process that launched the action.

   - **Block, kill and quarantine** SES Evolution blocks malicious actions, shuts down the process that launched the action and quarantines suspicious files. See Managing file quarantine.

8. In the upper banner in the rule, you can:

   - Select the version of protection to run, either a specific version, or **Always use the latest version**. If you are not using the latest version available, the indicator 🔼 will appear to the right of the drop-down list. It also appears on a policy's general panel when at least one enabled advanced protection mode is not using its latest version.

   - Enter a comment.

   - Select the log settings that this rule will send.

   - Specify whether an action must be performed when a log is sent for this rule.

9. Rules against **WMI persistence**, **Malicious use of certutil**, **Environment discovery**, **Ransomware** and **Parent PID Spoofing** each have specific parameters:

WMI Persistence	Compatibility list: in this section, list the consumers that represent legitimate WMI events and which the protection mode must not block.
Malicious use of certutil	Compatibility list: add here the IDs of applications which could use *certutil.exe* in legitimate cases and which should not be blocked by the protection.
Environment discovery	• Interval: indicate the interval in seconds (minimum five seconds) between the first command and the last command, and the interval after which discovery operations must be ignored. • Compatibility list: in this section, add the IDs of applications allowed to run commands similar to discovery operations and which the protection mode must not block. • Sensitivity: select the threshold above which the protection will be triggered.
Ransomware	• Compatibility list: add here identifiers for legitimate encryption applications that should not be blocked by protection, such as *StormshieldData Security*. • Sensitivity: select the threshold above which the protection will be triggered. With a **Very low** level, the protection mode will be triggered if a ransomware program encrypted at least 20 files within 3 seconds. With a **Low** level, the threshold will be 15 files, and with the **Moderate** level, 10 files. If you enable this anti-ransomware protection mode, ensure that you also Enabling Windows shadow copies so that you can restore lost files if necessary. For further information on restoration, refer to Managing ransomware attacks.
Protection against various threats	Compatibility: here, add the IDs of applications that will be allowed to spoof parent processes without being blocked by the protection mechanism.

10. Click on **Save** at the top right of the window to save changes.

> **ℹ NOTE**
> If subsequently, you want to change the version of an advanced protection, a deployment is required after the change is made.

8.5 Defining access control rules

To protect hosts and resources, SES Evolution makes it possible to control access to the registry base, files, processes, networks, volumes, devices and Wi-Fi access points. To do so, create security rule sets that will allow you to control access to these resources and build a security policy.

For every rule, you can define:

- How all applications behave by default with the resource targeted in the rule,
- Specific behavior for certain applications.

For more information on application behavior, refer to Using default behavior and specific behavior in rules.

The sequence of rule in a policy matters, because as soon as a rule matches a packet, the rules that are placed after this rule may not necessarily be read. The most specific rules must therefore be placed before more general rules. For more information on the sequence of rules, refer to Organizing rules and rule sets in a policy.

Access control rules can be created in the **Security > Policies** menu in the console, under the **Application**, **ACL resources**, **Networks** and **Devices** tabs in rule sets.

Most access control rules function in the same way:

- In the left section of the rule, define the resources that you want the rule to cover,
- In the right section, define the actors in the rule (specific behavior) and grant or deny them access privileges to the targeted resources. The actions that can be performed on various resources are different for each rule type, depending on whether you are in a protection rule set or an audit rule set. In audit rule sets, each action can be set to **Allow** and **Audit**.

In both cases, resources and actors are represented by the identifiers that must be created beforehand or created directly in the rule for some types of rules. For more information, refer to the section Creating identifiers.

Security rules can be disabled at any time. For more information, refer to the section Disabling security rules.

8.5.1 Controlling process creation

Malicious programs can strike by creating their own processes or creating them through third-party applications.

SES Evolution enables protection from such attacks.

Requirements

An application identifier must be created beforehand for the processes to be protected and for legitimate processes allowed to create other processes. For more information, refer to the section Creating application identifiers.

Creating a process creation rule

1. Select the **Security > Policies** menu and click on your policy.
2. Select a rule set.
3. Click on the **Application > Process creation** tab.
4. If you are in read-only mode, click on **Edit** in the upper banner.
5. Click on **Add > Rule (Process creation)**.
   A new line is displayed.
6. Click on ⎓＋ in the application ID area and select the process(es) to protect.
7. Click on the ≠＋ icon to select the process(es) to be excluded from protection.
8. In the **Creation** field of the **Default behavior** area, select a behavior from those available for this rule type:
   - **Allow** to allow the action by default,
   - **Block** to block the action by default,
   - **Block and kill** to block the action by default, and shut down the process that launched the action.
   - **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section [Managing file quarantine](#).
   - **Ask** for the user to be consulted.
   - **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
   - **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
   - **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
   - **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.
9. Click on **+ Add specific behavior** and select the process(es) that you want to exclude from the default behavior. In the associated **Creation** field, choose the desired behavior.

10. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⠿ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

  Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

11. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see **Classifying attacks according to the MITRE repository**.

12. Click on **Save** at the top right of the window to save changes.

> ✏️ **EXAMPLE**
> You can restrict the creation of the *rundll32* process only to Microsoft applications. In this case, select *rundll32* from the processes to be protected, select **Block** as the default behavior, then allow Microsoft applications in the specific behavior.

8.5.2 Controlling code execution

This protection type allows or prohibits the loading of executable code from executable files or DLL libraries.

The files or libraries in question are identified in rules by a path, alternate data stream, owner and/or volume type.

> ✏️ **EXAMPLE**
> These rules make it possible for example to allow the execution of only binary files installed by the operating system or by administrators of the agent pool, or to prevent dangerous applications from executing certain DLL files.

Requirements

An application identifier must be created beforehand for applications that are allowed or not allowed to run files or libraries. For more information, refer to the section **Creating application identifiers**.

Creating a code execution rule

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **Application > Code execution** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Code execution)**. A new line is displayed.

6. Click on `= +` in the area on the left to show the window where IDs of restricted access executable or DLL files are created.
   - And/or -

   Click on the `≠ +` icon to display the window for creating the ID of the executable file(s) or DLL(s) you wish to exclude from access control.

7. Enter the ID name.

8. Enter a path, an extension or name of an executable or DLL file. The generic characters "?" and "*" are allowed in this field.

9. Choose the type of volume on which the file or DLL is located.

10. You can specify the Windows account that owns the files in advanced settings, provided that these files are located on a local volume. You can also manually enter a Security ID (SID) to indicate a personal Windows account. This option makes it possible to allow or prevent the execution of files or DLLs hosted on certain accounts.

11. You can also specify an alternate data stream. A file's alternate data stream contains metadata and makes it possible to find out the origin of the file. For example, by specifying the alternate data stream "zone.identifier", rules can be created for files originating from the Internet. The alternate data stream can also be an attack vector by harboring malicious code. The generic characters "?" and "*" are allowed in this field.

12. Click on **OK** to close the ID creation window. Scroll over the name of the ID to see a summary of the settings.

> ✎ **EXAMPLES**
> - Prevent **lxssmanager.dll* from running on all applications.
> - Prevent **system.management.automation.dll* from running on all applications except legitimate applications.

13. In the **Execution** field of the **Default behavior** area, select a behavior from those available for this rule type: :
    - **Allow** to allow the action by default,
    - **Block** to block the action by default,
    - **Block and kill** to block the action by default, and shut down the process that launched the action.
    - **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.
    - **Ask** for the user to be consulted.
    - **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
    - **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
    - **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
    - **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

14. Click on **+ Add specific behavior** and choose the resource(s) that you want to exclude from the default behavior. In the associated **Execution** field, select the desired behavior.

15. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on [⠿] when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

  - Unclassified: unclassified rule.

  - Nominal: non-blocking rule conforming to nominal application behavior.

  - Protect: blocking rule with a high log severity level.

  - Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

  - Detect: non-blocking audit rule or passive rule.

  - Context: rule used to build an attack graph.

  - Syslog: rule triggering logs sent exclusively to a Syslog server.

  - Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

  Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

16. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

17. Click on **Save** at the top right of the window to save changes.

8.5.3 Controlling access to processes

Malicious programs can strike by accessing legitimate processes to retrieve sensitive data or inject malicious code into them.

Rules that regulate access to SES Evolution processes enable protection against such attacks without completely blocking inter-process communication, some of which is legitimate.

Access to a process or thread cannot be fully blocked, but you can restrict privileges during this operation.

These rules only apply to applications. They do not apply to drivers.

> ✎ **EXAMPLE**
> Blocking access to the memory of a process can prevent passwords from being stolen from a browser's memory when it is open.
> Other examples are given at the end of this section.

Requirements

An application identifier must be created beforehand for the processes to be protected and for legitimate processes allowed to access other processes. For more information, refer to the section Creating application identifiers.

Creating a rule for access to processes

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **Application > Process access** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Process access)**.
   A new line is displayed.

6. Click on ⊟ ⊞ in the application ID area and select the process(es) to protect.

7. Click on the ≠ ⊞ icon and select the process(es) to be excluded from protection.

8. In the **Default behavior** field, select the behavior for each action (in audit rule sets, only the **Read** action can be configured):

- **Read**: choose what the rule must do when it reads the memory of the process.
- **Modification**: choose the rule action in the event of process memory modification.
- **Execution flow tampering**: a program that takes control of a process can modify its execution pointer. Choose what the rule must do when the execution flow of the process is tampered with.
- **Handle duplication**: choose what the rule must do when a process attempts to duplicate a resource that belongs to another process.

The list of all behaviors is described below:

- **Allow** to allow the action by default,
- **Block** to block the action by default,
- **Block and kill** to block the action by default, and shut down the process that launched the action.
- **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.
- **Ask** for the user to be consulted.
- **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
- **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
- **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
- **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

9. Click on **+ Add specific behavior** and select the process(es) that you want to exclude from the default behavior. Select the behavior for each case.

10. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⊞ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- Unclassified: unclassified rule.

- Nominal: non-blocking rule conforming to nominal application behavior.

- Protect: blocking rule with a high log severity level.

- Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- Detect: non-blocking audit rule or passive rule.

- Context: rule used to build an attack graph.

- Syslog: rule triggering logs sent exclusively to a Syslog server.

- Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

  Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

11. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

12. Click on **Save** at the top right of the window to save changes.

✎ **EXAMPLES**

You can block any application from accessing the password manager to prevent hackers from accessing passwords or injecting code into its process. In this case, choose the password manager from the list of processes to be protected and select **Block** for all actions in the default behavior. Do not define any specific behavior.
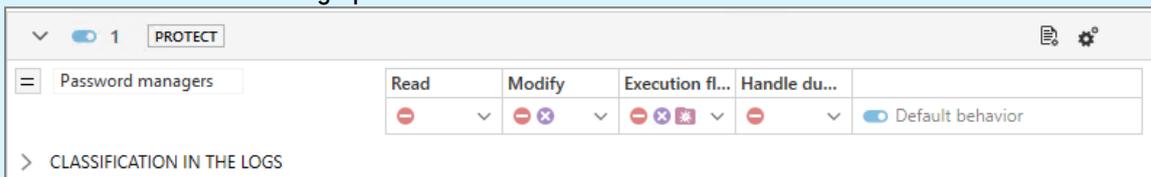
You can also block execution flow tampering for major applications such as business applications, to prevent hackers from shutting them down or suspending them.
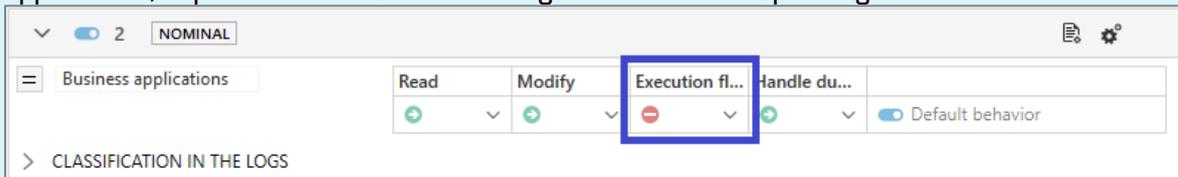
8.5.4 Protecting against code injection

Code can be injected into an application to make it run code from another application. SES Evolution makes it possible to protect your applications against the injection of malicious code.

✎ **EXAMPLE**

There are two possible approaches, illustrated as follows:

- Use case 1: No applications are allowed to inject code, except clearly identified legitimate applications (e.g., antivirus, Windows error reporting, etc.). This is the most commonly used approach.
- Use case 2: No applications are allowed to inject code in the password manager.

Requirements

An application identifier must be created beforehand for every application to be protected and for every application allowed to inject legitimate code. For more information, refer to the section Creating application identifiers.

Creating a rule to protect against code injection

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **Application > Code injection** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Code injection)**.
   A new line is displayed.

6. Click on the [= +] icon in the application ID area and select the application(s) affected by the default behavior.
   For use case 1, do not add any applications since you will be protecting all of them.
   For use case 2, add the password manager.

7. Click on the [≠ +] icon to select the process(es) to be excluded from protection.

8. In the **Access** field of the **Default behavior** area, select a behavior from those available for this rule type: :

   • **Allow** to allow the action by default,

   • **Block** to block the action by default,

   • **Block and kill** to block the action by default, and shut down the process that launched the action.

   • **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.

   • **Ask** for the user to be consulted.

   • **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.

   • **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.

   • **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.

   • **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

9. Click on **+ Add specific behavior** and select the application(s) that you want to exclude from the default behavior.
   For use case 1, add the applications that inject legitimate code (e.g., antivirus, Windows error reporting, etc.) and in **Access**, choose **Allow**.
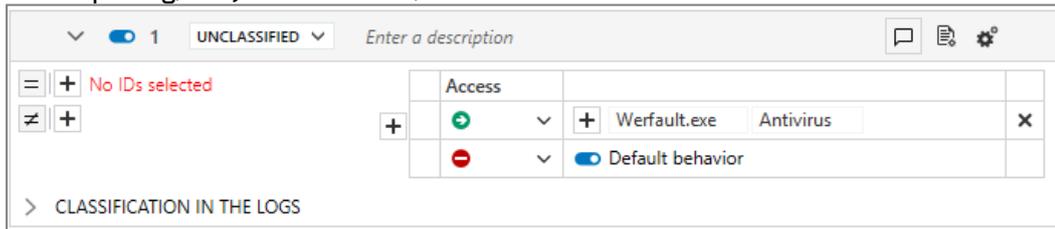
For use case 2, do not add any applications since the password manager will be fully
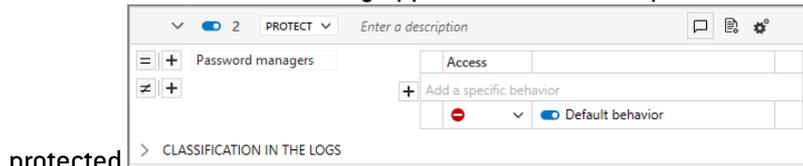
protected.

10. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⬚ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

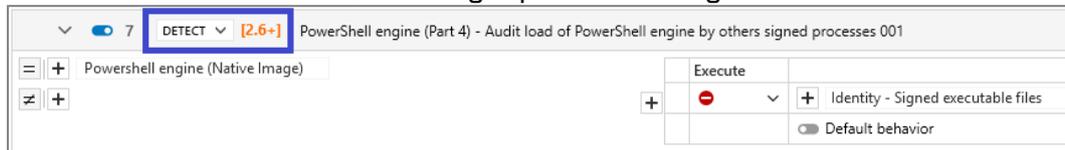 ○ Unclassified: unclassified rule.

  ○ Nominal: non-blocking rule conforming to nominal application behavior.

  ○ Protect: blocking rule with a high log severity level.

  ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

  ○ Detect: non-blocking audit rule or passive rule.

  ○ Context: rule used to build an attack graph.

  ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

  ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

  Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

11. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

12. Click on **Save** at the top right of the window to save changes.

8.5.5 Protection against keylogging

Keylogging enables hackers to capture all of a user's keystrokes in order to steal passwords, confidential data, etc. It acts on targeted applications.

SES Evolution prevents foreground applications from sending their keystrokes to other applications. However, it can receive its own keystrokes.

For more global protection against any use of the SetWindowsHookEx API, enable protection against application hooking instead. For more information, refer to Protection against various threats and Configuring threat protection.

> ✏️ **EXAMPLE**
> You can use this protection to block keylogging on web browsers, password managers, and the Windows file explorer. Allow them only for legitimate applications such as virtualization and remote control tools.

Requirements

An application identifier must be created beforehand for every application to be protected and for every application allowed to log keystrokes. For more information, refer to the section Creating application identifiers.

Creating a rule to protect against keylogging

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **Application > Keylogging** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Keylogger)**.
   A new line is displayed.

6. Click on ⊟ ⊞ in the application ID area and select the application(s) to protect.
   Add *Internet Explorer*, *Windows Explorer* and *password manager* for example.

7. Click on the ≠ ⊞ icon to select the application(s) to be excluded from protection.

8. In the **Status** field of the **Default behavior** area, choose a behavior from among those available for this rule type: :

- **Allow** to allow the action by default,
- **Block** to block the action by default,
- **Block and kill** to block the action by default, and shut down the process that launched the action.
- **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section **Managing file quarantine**.
- **Ask** for the user to be consulted.
- **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
- **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
- **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
- **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

9. Click on the **+ Add a specific behavior** icon and select the application(s) that you want to allow.
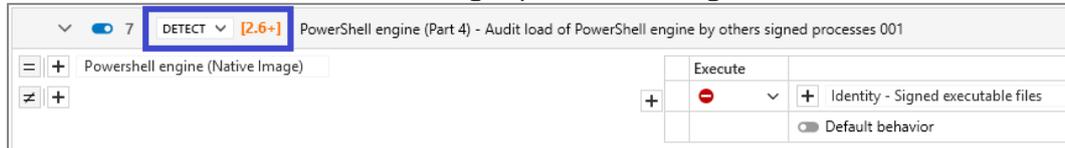Add applications that legitimately log keystrokes, e.g., *remote control tools*, and in the **Status**, select **Allow** so that the protection will allow these applications.

10. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⊞ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

  Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

11. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

12. Click on **Save** at the top right of the window to save changes.

8.5.6 Controlling access to files

This protection mode makes it possible to control specific applications' access to files. These files are identified in rules by a path, alternate data stream, owner and/or volume type.

> ✏️ **EXAMPLE**
> You can protect all your Microsoft Office files and other sensitive files so that they can be modified only by legitimate applications such as Windows Explorer, Office suite, Windows tools, etc. Other applications will be granted read-only access to these files.

Requirements

An application identifier must be created beforehand for applications that are allowed to access files and for those that you want to block. For more information, refer to the section Creating application identifiers.

Creating a file access rule

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **ACL resources > File** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Files)**.
   A new line is displayed.

6. Click on ⊟ ➕ in the area on the left to show the window where IDs of restricted access files are created.
   - And/or -

   Click on the ≠ ➕ icon to display the window for creating the identifier of the file(s) you wish to exclude from access control.

7. Enter the ID name.

8. Enter a file, path or extension. The generic characters "?" and "*" are allowed in this field. Full paths beginning with a letter (i.e., *E:\Data\Backup*) are not supported if the **Volume type** is remote or removable.
   Stormshield strongly recommends the use of EsaRoots path roots provided by SES Evolution instead of drive letters (i.e., *C:\...*). In fact, these letters may differ from one workstation to another.

   > ℹ️ **NOTE**
   > You can enter a path that contains the letter of a local hard disk or SSD drive in this field. However, if users change the letter of the drive or add one, you must restart the workstation or modify the policy that the agent applies so that the drive can be detected.

9. Choose the type of volume on which the file or file type is located.

10. You can specify the Windows account that owns the files in advanced settings, provided that these files are located on a local volume. You can also manually enter a Security ID (SID) to indicate a personal Windows account. This option makes it possible to allow or prevent access to files hosted on certain accounts.

11. You can also specify an alternate data stream. A file's alternate data stream contains metadata and makes it possible to find out the origin of the file. For example, by specifying the alternate data stream "zone.identifier", rules can be created for files originating from the Internet. The alternate data stream can also be an attack vector by harboring malicious code. The generic characters "?" and "*" are allowed in this field.

12. Click on **OK** to close the ID creation window. Scroll over the name of the ID to see a summary of the settings.

13. In the **Default behavior** field, select a behavior from among those available for this type of rule: :

    - **Allow** to allow the action by default,
    - **Block** to block the action by default,
    - **Block and kill** to block the action by default, and shut down the process that launched the action.
    - **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.
    - **Ask** for the user to be consulted.
    - **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
    - **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
    - **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
    - **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

14. Click on **+ Add specific behavior** and choose the resource(s) that you want to exclude from the default behavior. Select the behavior for each case.
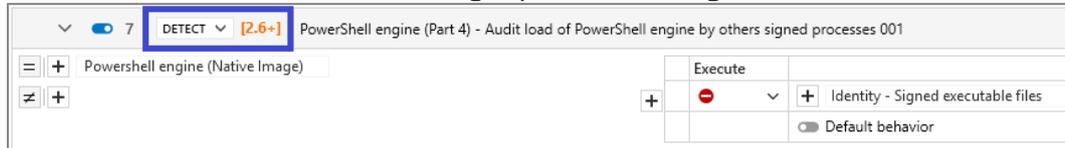
> ✎ **EXAMPLE**
> Block the ability to modify or delete Office files and other sensitive files by default. Allow these actions only for legitimate applications.
>
>

15. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⊞ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

16. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

17. Click on **Save** at the top right of the window to save changes.

8.5.7 Controlling access to the registry base

This protection type makes it possible to control specific applications' access to keys and values in the registry base. As such, access to particularly sensitive keys can be protected, as they are a prime target of malicious programs.

> ✏️ **EXAMPLE**
> To prevent a malware program from disabling Windows security tools via the registry base, you can protect their registry keys so that they can only be modified by legitimate Windows applications.

Every registry path can be a full path or contain the generic characters "?" and "*".

Requirements

Application identifiers must be created beforehand for applications that are allowed to access registry and for those that you want to block. For more information, refer to the section Creating application identifiers.

Creating a registry access rule

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **ACL resources > Registry** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Registry)**.
   A new line is displayed.

6. In the left-hand area, click on the ⊟➕ icon to display the window for creating the identifier of the registry keys whose access you wish to control.
   - And/or -

   Click on the ≠➕ icon to display the window for creating the identifier of registry keys you wish to exclude from access control.

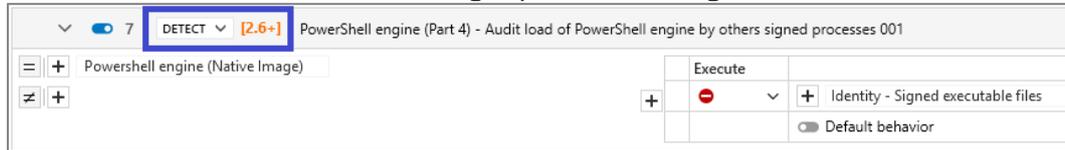7. Enter the ID name.

8. Enter the path to the key.

> 💡 **TIP**
> The path to the key can be copied from the registry base and pasted in the **Key** field.

9. Choose where to apply these rules:

- **Key and Values**. These rules cater to the most frequent protection needs. If you do not enter a value, all the values of the key will be protected, including the key itself. If you enter a single value, the other values of the key will not be protected.
- **Key**: These rules provide more advanced protection. Only the key is protected, but its values are not.
- **Values**: These rules also provide more advanced protection. Only the values are protected, but the rule does not protect the key itself. Even if the values of a key are protected from deletion, if the deletion of the key itself is allowed, the values may be deleted together with the key.

10. Click on **OK** to close the ID creation window. Scroll over the name of the ID to see a summary of the settings.

11. In the **Default behavior** field, select a behavior from those available for this type of rule: :

- **Allow** to allow the action by default,
- **Block** to block the action by default,
- **Block and kill** to block the action by default, and shut down the process that launched the action.
- **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.
- **Ask** for the user to be consulted.
- **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
- **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
- **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
- **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

12. Click on **+ Add specific behavior** and choose the resource(s) that you want to exclude from the default behavior. Select the behavior for each case.

> ✏️ **EXAMPLE**
> By default, block access to the registry keys of Windows security tools such as Windows Defender, Windows Firewall, etc. Allow only legitimate processes to perform these operations, e.g., Windows update and software installer, security solutions, etc.

13. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on [icon] when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

  Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

14. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

15. Click on **Save** at the top right of the window to save changes.

8.5.8 Controlling access to the volume

This protection prevents applications from bypassing security checks that the file system of the system disk conducts, and makes it possible to access the raw volume directly.

In the rules, you can allow or prohibit access to the raw volume by the applications of your choice.

In whitelist mode, a single rule may be enough to authorize access for some applications and block it for all others. You must create several rules if you want to select different log settings. In this case, define "Block" as the default behavior in only the last rule.

> **✎ EXAMPLE**
> Example of a rule prohibiting all applications, except legitimate applications, from accessing the volume.
>
>

Requirements

An application identifier must be created beforehand for applications that are allowed or not allowed to access the raw volume. For more information, refer to the section Creating application identifiers.

Creating a volume access rule

1. Select the **Security > Policies** menu and click on your policy.
2. Select a rule set.
3. Click on the **ACL resources > Volume** tab.
4. If you are in read-only mode, click on **Edit** in the upper banner.
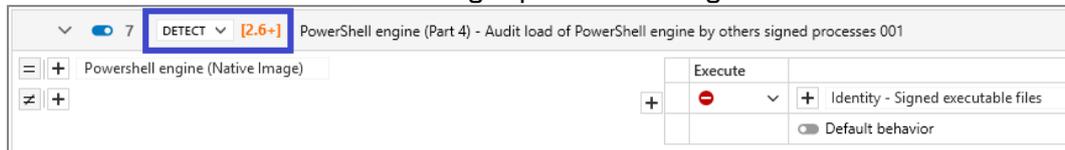5. Click on **Add >Rule (Volume)**. A new line is displayed.

6. In the **Access** field in the **Default behavior** area of a protection rule, select the behavior that applies to all applications that may access the raw volume:

   - **Allow** to allow the action by default,
   - **Block** to block the action by default,
   - **Block and kill** to block the action by default, and shut down the process that launched the action.
   - **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.
   - **Ask** for the user to be consulted.
   - **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.
   - **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.
   - **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.
   - **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

7. Click on **+ Add a specific behavior** and choose the resource(s) that you want to exclude from the default behavior. In the associated **Access** field, select the desired behavior.

8. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⊞ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

9. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

10. Click on **Save** at the top right of the window to save changes.

8.5.9 Controlling network access

This protection is used to control access to incoming and outgoing networks by specific applications.

Access can be filtered by:

- Network events such as "bind", "accept" (server rule) and "connect" (client rule),
- TCP and UDP protocols,
- Specific ports,
- Specific IPv4 or IPv6 addresses.

It is not necessary to explicitly open communications between the SES Evolution server and the agents. Indeed, the agent's self-protection mechanism ensures that no security rule whatsoever can block these communications.

> ✏️ **EXAMPLE**
> Network rules make it possible to:
>
> - Protect a server by controlling access to the host,
> - Force users of a service in the company to use a specific application to access a given network resource.

Requirements

The following must be created beforehand:

- Application IDs for allowed applications or applications that cannot access the network. For more information, refer to the section Creating application identifiers.
- Network IDs for the IP addresses that you want to protect. For more information, refer to the section Creating network identifiers.

Creating a network access rule

There are two types of rules; client rules and server rules.

- As part of a rule set that applies to workstations, client rules allow or do not allow applications to connect to remote resources (**Remote** field) by controlling the "connect" network event. They also make it possible to cater to specific subnets for example (**Local** field).
- As part of a rule set that applies to servers, server rules allow or do not allow applications to open ports and accept incoming connections (**Local** field) by controlling the "bind" and "accept" network events. They also make it possible to specify the source of connections (**Remote** field).

To create a network access rule:

1. Select the **Security > Policies** menu and click on your policy.
2. Select a rule set.

3. Click on the **Networks > firewall** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Choose to add a Client network rule or a Server network rule by clicking one of the **Add** buttons. A new line is displayed.

6. Choose the network IDs of the resources you want to protect in the left side of the rule:

   - **Local** $\boxed{= \ +}$: Local resource affected by the rule. E.g., if the workstation has several network cards, you can specify which card is impacted.

   - **ELocal** $\boxed{\neq \ +}$ Local resource excluded from the rule.

   - **Remote** $\boxed{= \ +}$: Remote resource affected by the rule. E.g., the internet.

   - **Remote** $\boxed{\neq \ +}$: Remote resource excluded from the rule.

7. In the **Ports** field, indicate the ports affected by the network rule. These ports are the destination ports for client rules and local ports for server rules.

   - To add several ports at one go, separate them with commas. Example: 8080.8081.

   - To add a port range, separate the first value and last value with a dash. Example: 80-90

   - Leave the field empty to specify that all ports are concerned.

8. Choose the TCP or UDP transport protocol, or both.

9. In **Default behavior**, choose the behavior of each Connect, Accept or Bind network event:

   - **Accept** (for server rules): allows or does not allow specified applications to receive incoming connections on the network resource(s) indicated,

   - **Bind** (for server rules): allows or does not allow specified applications to open connections on the network resource(s) indicated,

   - **Connect** (for client rules): allows or does not allow specified applications to connect to the network resource(s) indicated.

   Protection rules can behave as follows:

   - **Allow** to allow the action by default,

   - **Block** to block the action by default,

   - **Block and kill** to block the action by default, and shut down the process that launched the action.

   - **Block, kill and quarantine** to block the action by default, kill the process that triggered the action, and quarantine suspicious files. For more information, see the section Managing file quarantine.

   - **Ask** for the user to be consulted.

   - **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.

   - **Skip rule** to ignore the rule contained in this rule set and evaluate the next rule.

   - **Skip rule group** to ignore the rules contained in the rule group and evaluate the next rule group or rule.

   - **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

10. Click on **+ Add a specific behavior** and choose the application identifiers of resource(s) that you want to exclude from the default behavior.

✏️ **EXAMPLE**
This is the client rule that can **block** connections from the network card on the unprotected

---

network card to the internet and the protected network over ports 80, 443 and 8080 and TCP. Only the web server specified in the protected network can be accessed.

11. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⠿ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

12. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

13. Click on **Save** at the top right of the window to save changes.

8.5.10 Controlling Wi-Fi access

This protection mode controls how mobile workstations access Wi-Fi networks by:

- Allowing or preventing the use of Wi-Fi connections and defining a whitelist of Wi-Fi access points in the form of rules, based on the SSID of the Wi-Fi network and/or MAC address of the Wi-Fi access point,
- Allowing or preventing the use of ad hoc Wi-Fi connections,
- Forcing the use of secure authentication protocols.

Wi-Fi connections are disabled by default in protection rule sets. If there are several protection rule sets in your security policy, ensure that you enable the policy only for the set(s) in which you want to configure Wi-Fi access, and arrange your rule sets in the right order in the policy. If you enable and allow Wi-Fi access in a rule set near the top of the policy, this rule may overload and cancel the effect of the Wi-Fi access configuration in the rule sets that follow.

Depending on certain events, the block policy for Wi-Fi connections inside or outside a perimeter can be enabled using conditional policies. For more information, refer to the section Assigning a security policy to agents.

Allowing or blocking Wi-Fi connections

To allow or block the Wi-Fi connection feature on workstations:

1. Select the **Security > Policies** menu and click on your policy.
2. Select a rule set.
3. Click on the **Networks > Wi-Fi** tab.
4. If you are in read-only mode, click on **Edit** in the upper banner.
5. The first **Wi-Fi connection** rule cannot be deleted, and is disabled by default. This rule allows you to authorize or block the operation of WiFi network cards on workstations. It is only present in a protection rule set. Expand the rule, enable it by clicking on the ⬜ button on the left, then authorize or block access.
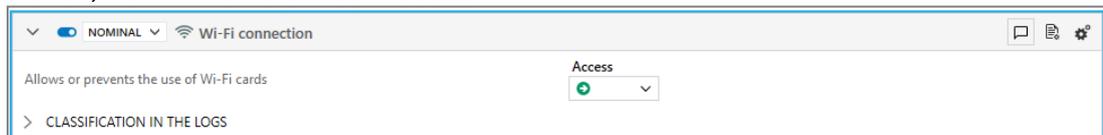
If you disable or block Wi-Fi access and your policy contains rules regarding access to Wi-Fi networks, these rules will not be scanned.

For more granular management of access to Wi-Fi networks, allow Wi-Fi connections and create **Wi-Fi network** rules.

Controlling access to Wi-Fi networks

After you allow the Wi-Fi connection in the first rule of a protection rule set, create rules to block or allow access to certain Wi-Fi networks, or create rules to audit access to Wi-Fi in an audit rule set. By default, if no rules are defined, access to all Wi-Fi networks is allowed and rules can therefore be used to block access to networks in blacklist mode. If you prefer to operate in

whitelist mode, i.e., explicitly allowing access to certain networks, create a rule that blocks access to all networks other than those allowed, and place this rule at the end.

To create WiFi network rules:

1. In the **WiFi** tab, click on **Add > Rule (WiFi networks)**. A new line is displayed.

2. In the left side of the rule, click on [= || +] to add a Wi-Fi network.

3. Enter the following information:

    -   Network name,
    -   SSID (Service Set IDentifier). The use of wildcards is permitted (e.g.: *stormshield**) and case is not important,
    -   MAC access of the access point(s) in hexadecimal format. To indicate several, click on the + icon,
    -   WiFi connection mode,
    -   Authentication type, to secure communications with the Wi-Fi access point(s).

    > **ℹ NOTE**
    > The WPA3 authentication mode is not compatible with SES Evolution agents in a version lower than version 2.4.

4. In the **Connection** field, select **Allow** or **Block**.

5. In the upper banner in the rule, you can:

    -   Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
        Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.
    -   Select the log settings that this rule will send.
    -   Specify whether an action must be performed when a log is sent for this rule.
    -   Enter a comment.
    -   Enter a description to explain what this rule aims to achieve.

6. The row number of each rule appears on its left. Rearrange the sequence of your rules if you need to, by clicking on the arrows above and below the row number.

7. Click on **Save** at the top right of the window to save changes.

8.5.11 Allowing temporary web access

The temporary web access mechanism allows a user to bypass Network protection rules in the policy, with specific applications and for a duration that you can set.

When this duration expires, new connections will be blocked once more, according to the rules in the security policy. However, applications for which connections were opened during temporary web access will not be shut down, and existing connections will not be interrupted.

> **✎ EXAMPLE**
> Temporary web access makes it possible to manage mobile users who want to log in to their corporate network via a VPN tunnel from unsecure networks. When these workstations are

> outside the corporate network, the security policy that applies may prevent communications over the network. Temporary web access therefore allows them to temporarily unblock the VPN client and browser upon users' request, so that the client can log in to the corporate network and switch to the internal security policy. Users will then be able to use their workstations normally.

Temporary access only needs to be allowed on one of the policies assigned to an agent group for this feature to be available on the agent side.

The temporary web access feature is available only in protection rule sets.

This feature is disabled by default. If there are several protection rule sets in your security policy, ensure that you enable the policy only for the set(s) in which you want to configure temporary web access, and arrange your rule sets in the right order in the policy. If you enable and allow temporary web access in a rule set near the top of the policy, this rule may overload and cancel the effect of the temporary web access configuration in the rule sets that follow.

Prerequisites

- Application identifiers must be created beforehand for applications allowed to access unrestricted networks when temporary web access is enabled. For more information, refer to the section Creating application identifiers.
- The application firewall must be enabled in the configuration of agent groups:
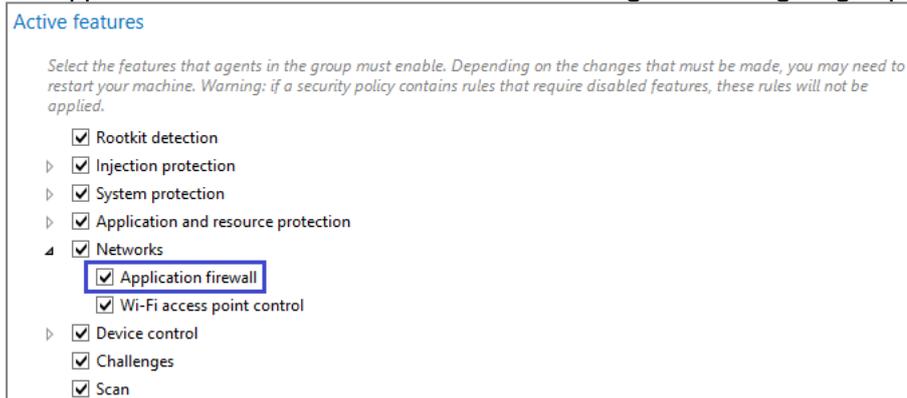
Allowing temporary web access

1. Select the **Security > Policies** menu and click on a policy.
2. Select a protection rule set.
3. Click on the **Networks > Temporary web access** tab.
4. If you are in read-only mode, click on **Edit** in the upper banner.
5. Enable the feature.
6. Allow temporary web access.
7. Select one or several application identifiers allowed to access the web. These applications will be allowed to connect to all IP addresses over all ports.
8. Choose the maximum duration of web access.
9. Choose the number of access authorizations. The counter will be reset when the user restarts the workstation.
10. If necessary, create a shortcut on the user's desktop. Users have several ways to enable temporary access on their workstations. For more information, see the next section.
11. Click on **Save** at the top right of the window to save changes.

Accessing the web temporarily from the agent

The SES Evolution agent provides the user with several ways to enable temporary web access:

- A **Start temporary web access** button in the [icon] tab of the agent's interface,

Remaining web access	3
⊕ Start temporary web access	

- A pop-up menu that appears by right-clicking on the agent's [icon] icon in the taskbar.
- A desktop icon, if the feature is enabled in the temporary web access settings,
- The command `EsGui /GrantWebAccess` to be inserted into a script, for example.

When the user's temporary web access is in progress, a banner at the bottom of the agent's interface indicates the remaining time.

The user can stop temporary access:

- via the agent's interface,
- via the pop-up menu of the agent's icon in the taskbar.

8.5.12 Controlling access to devices

SES Evolution allows you to control access to all types of devices that can be connected to users' workstations.

Controlling access to general devices

Controlling access to Bluetooth devices

Controlling access to USB devices

Controlling storage on USB devices

Controlling application execution from removable devices

8.6 Importing Sigma security rules

The Sigma format is a standard unified language for describing log-based incident detection rules. In particular, Sigma rules can be used to create and share standardized detection rules that can be used regardless of the SIEM or system.

They are written in a text file in YAML format and sent to SES Evolution via its public API by a SIEM or a security administrator. For more information, see the Sigma Documentation.

Stormshield has developed two import scripts used to send Sigma rules to the SES Evolution API, convert them to SES Evolution rules, and deploy them to the agents.

If you want to use API requests directly without the help of Stormshield scripts, see Enabling and managing SES Evolution's public API and the API documentation.

You can enable Sigma security rules via advanced protection: *Sigma advanced protection*.

8.6.1 Requirements

- For security reasons, use only official scripts provided and signed by Stormshield.
- The SES Evolution root certification authority must be present on the machine where the Sigma rule import scripts are running. This ensures that the communication between the backend and the workstation is trusted and that the API key will never be passed on to a third party. For more information, see Importing the SES Evolution root certification authority.

> **ⓘ NOTE**
> Machines hosting SES Evolution back office components already have the required certification authority.

- Sigma rules imported into SES Evolution must be *detection* rules of the *Windows log event* or *Windows process creation* type, or *filter* rules.
- You must have an API key with *Policy* usage. For more information, see Adding an API key. Make sure you only share this key with people you trust.
- For security reasons, in Windows, the script requires a version of the PSReadline module higher than 2.0.4. For more information, see the Microsoft Documentation.
- Under Linux, the import script requires:
  - Installation of OpenSSL v3.x or higher.
  - Installation and activation of the virtual environment for Python: `pip` and `venv` tools. For more information, see the Python documentation.
  - Installation of the script prerequisites, whose *requirements.txt* file can be downloaded from your MyStormshield customer area, in the **Downloads > SES Evolution > Tools** section.

8.6.2 Importing the SES Evolution root certification authority

You must export the root certification authority (CA), then import it to the machine on which you are going to run the scripts Stormshield.

Exporting the CA

1. Go to a machine that hosts a SES Evolution backoffice component: backend, agent handler, or administration console.
2. In the Certificate Manager, select the *SES Evolution Root CA* and export it.
3. Copy the CA to the machine where you want to import it.

Importing the CA on a Windows machine

- Double-click the CA file and follow the instructions in the certificate import utility.

Importing the CA onto a Linux machine

1. Install the `ca-certificates` tool.
2. Convert the previously exported CA to *.crt* format using the following command:
   ```
   openssl x509 -inform DER -in sesrootca.cer -out sesrootca.crt
   ```

3. Copy the file using the following commands:

Ubuntu/Debian	Red Hat
`cp sesrootca.crt /usr/local/share/ca-certificates`	`cp sesrootca.crt /etc/pki/ca-trust/source/whitelist/`

4. Finalize the CA import using the following commands:

Ubuntu/Debian	Red Hat
`update-ca-certificates`	`update-ca-trust`

8.6.3 Importing and deploying the Sigma rules via the Stormshield script

1. In your MyStormshield client area, **Downloads > SES Evolution > Tools** section, download the *insert-sigma-rules-from-folder.ps1* scripts for Windows, and *insert-sigma-rules-from-folder.py* scripts for Linux.

2. Group all your Sigma rules together in the same directory, including linked filter rules.

3. On a Windows computer, open a PowerShell window in the directory where the scripts are located and run the following command:
   ```
   .\insert-sigma-rules-from-folder.ps1 -Directory Sigma_Rules_
   Directory -HostName Backend_Address -ApiKey API_Key********
   ```

   On a Linux computer, run the following command:
   ```
   python3 ./insert-sigma-rules-from-folder.py -directory Sigma_Rules_
   Directory -hostname Backend_Address -apiKey API_Key********
   ```

SES Evolution imports the Sigma rules, converts them, and deploys them to the agents.

The import status of each rule is shown on the screen, with an error message if a rule import fails.

A new version of *Sigma Advanced Protection* is created, visible in the administration console. Its description displays the number of rules it contains and its update date in UTC format. For further information, see Configuring advanced protections.

Individual rules are not visible in the console. To view the list of rules, use the API request */sigma/last-import-state*. For further information, see the API documentation.

> ⚠️ **CAUTION**
> For security reasons, after using the script under Linux, make sure that the terminal history no longer contains any authentication information. Use the `history -d <line number>` command to delete the lines concerned.

8.6.4 Enabling Sigma rules in a security policy

There are two methods for activating Sigma rules in SES Evolution.

Use the *Stormshield – Sigma Protection* rule set

1. In your MyStormshield client area, **Downloads > SES Evolution > Tools** section, download the *Stormshield – Sigma protection.cab* shared rule set.

2. Add the rule set to the desired policies. For more information, see Creating a security policy

By default, the rule set uses the latest version of Sigma Advanced Protection.

Enabling Sigma Advanced Protection in an existing rule set

For further information, see Configuring advanced protections.

8.6.5 Knowing the specifics of Sigma logs

Like other logs, Sigma logs are visible in the agent interface and in the **Agent logs** panel of the administration console. However, they have the following specificities:

- When Sigma rules are triggered by a process or event in the Windows log, they output a *Protection log.*
- The rule navigation button is hidden on these logs,
- Information such as the rule title, author, and ID is displayed in the log details,
- SES Evolution converts the severity levels of Sigma logs as follows:

Sigma	SES
Informational	Notice
Low	Warning
Medium	Error
High	Critical
Critical	Alert

For more information, see Monitoring SES Evolution agent activity.

8.7 Grouping security rules

If your rule sets contain a large number of rules, reading and maintaining them can prove difficult. In this case, you can create groups containing all similarly themed rules. For example, group together all the rules concerning Microsoft Office applications.

8.7.1 Creating a rule group

1. In the rules panel, click on the **Edit** button in the top banner.
2. Select one or more rules, then click **Add > Group from selection** or use the shortcut **CTRL + G**. The rules are combined into a single group.
3. In the rule group header, enter a name or description to identify the group, for example *Office applications*.
4. To the right of the group header, click on the color palette to select the group color.
5. To add or remove rules from the group, or to order rules within the group, use the drag-and-drop icon.

> **ⓘ NOTE**
> You cannot create new rule groups for built-in threat protection rules.

8.7.2 Disabling a rule group

1. In the rules panel, click on the **Edit** button in the top banner.

2. In the rule group header, disable the [switch icon] switch. The group is grayed out. When a group is disabled, all the rules it contains are inactive.

8.7.3 Deleting a rule group

1. In the rules panel, click on the **Edit** button in the top banner.
2. In the rules group header on the right, click on the **Delete** icon.
   The **Delete group** window appears.
3. Select the degree of deletion:
   - **Delete group only**: the rules are removed from the group but not deleted.
   - **Delete group and rules**: the rules are deleted along with the group.

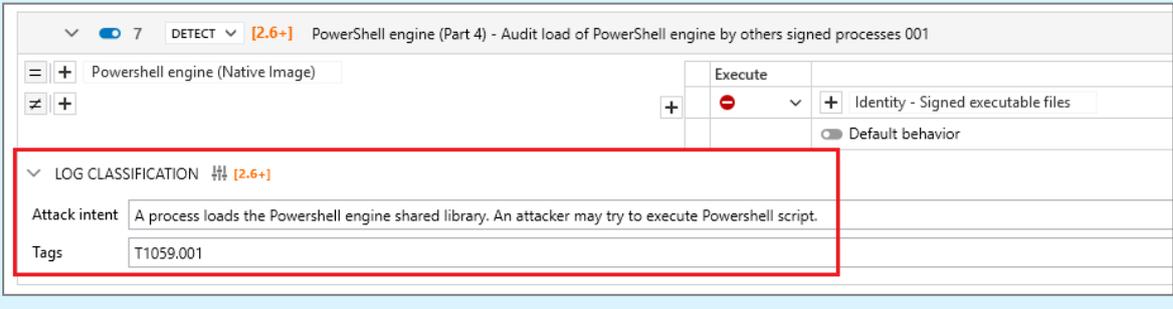## 8.8 Classifying attacks according to the MITRE repository

SES Evolution references in its logs the techniques and sub-techniques of attackers such as those listed and described in the MITRE ATT&CK® matrices and in the common vulnerabilities and exposures (CVE) published by the US organization MITRE. Thus, in the event of an attack, IT administrators can quickly identify it and take appropriate action.

This feature associates a security rule with an attack intent and tags.

This allows you to specify for each rule the type of attack that might be underway when the rule is applied. You can also specify a list of tags to automatically associate your SES Evolution rules with the MITRE repository.

The purpose of this feature is to rapidly provide administrators with information via the logs sent to the console when security rules are applied. They can then identify the possible attack underway on the pool by viewing its classification, and go directly to the URL of the MITRE technique or CVE. Tags can also reference vulnerabilities identified by Stormshield on its https://advisories.stormshield.eu/ website.

8.8.1 Adding intent and tags to a security rule

In each security rule, you can add an attack intent and tags corresponding to the techniques, sub-techniques or vulnerabilities referenced by MITRE. This information can be used to associate rules with known types of attack. They are displayed in logs when a rule is applied,

making it possible to classify logs generated by agents and quickly identify the type of attack possibly underway.

In a security rule:

1. Expand the **Classification in logs** part.

2. Indicate the intent of the attack, i.e. specify the type of attack that may be underway if the rule is triggered. For example, for a rule protecting a USB stick in read mode, you could specify "Extract sensitive data via removable media".

3. You can specify up to 10 tags, referring to the MITRE ATT&CK® matrices, CVE vulnerabilities or Stormshield. For tag formats, see MITRE ATT&CK® matrices, common vulnerabilities and exposures (CVE) or the https://advisories.stormshield.eu/ site. For example T1546.001, CVE-2021-40444 or STORM-2023-022.

You can use the search field to filter rules by attack intent or tags.

8.8.2 Viewing intents and tags in logs

Classification information is displayed differently on the SES Evolution agent and in the administration console.

The **Events** tab of the agent interface lets you view the tags associated with the rule that generated a log in the **Raw log** display:

In the agent logs in the administration console, the **Classification** panel in a log's details displays the attack intent and the tags associated with the rule that generated the log:

- Click on the MITRE and CVE tags to go directly to the details.

8.9 Defining rules for external events

External event audit rules allow you to collect certain events that occur on workstations, but which did not originate from standard SES Evolution components:

- Windows events,
- Events that the OSSEC analysis engine reported.

When the rule is enabled, collected external events will appear as logs in the **Agent logs** panel of the administration console and on the SES Evolution agent interface.

8.9.1 Forwarding Windows events in SES Evolution

The forwarding of Windows events consists of indicating in a rule which logs and which Windows events SES Evolution must collect and display.

> ### ✐ EXAMPLE
> You can choose to forward events relating to user connections on workstations, to monitor who logged in and when.

Create an event forwarding rule:

1. Select the **Security > Policies** menu and click on your policy.
2. Select an audit rule set.
3. Click on **External events > Event forwarding**.
4. If you are in read-only mode, click on **Edit** in the upper banner.
5. Click on **Add > Rule (Event forwarding)**.
   A new line is displayed.

6. Click on **+ Monitored events** and provide the following information:
Log name
Enter the name of the Windows log, e.g., *Security*, *Microsoft-Windows-Windows Defender/Operational*. To find out the name of a log, look up its properties in the Windows Event Viewer.
You can monitor a log that is not enabled in Windows. In this case, SES Evolution will automatically enable it. However, keep in mind that if there are too many events in this log, it may impact the performance of Windows.
If you enter a filter request in XML in the next field, the **Log name** is not completely necessary.

Filter request
If needed, enter a filter request to collect only some events in the log. To obtain a request:

1. Open the Windows Event Viewer.

2. Right-click on the log of your choice > **Filter the current log**.

3. In the **Filter** tab, select your filtering options.

4. Copy the contents of the **XML** tab and paste it in the **Filter request** field in the window of the event forwarding rule.

You can also manually enter a request in XPath. Enter for example the log name `Security` and the filter request `*[System[(EventID=4625)]]` to retrieve all events with the ID 4625 in the Security log.

Description
Enter a description if necessary.

You can also import a custom Windows events view, which will automatically fill in all fields with the desired values. To do so, go to the Windows Event Viewer and export the desired custom view in XML, and import it by clicking on the arrow on the right.
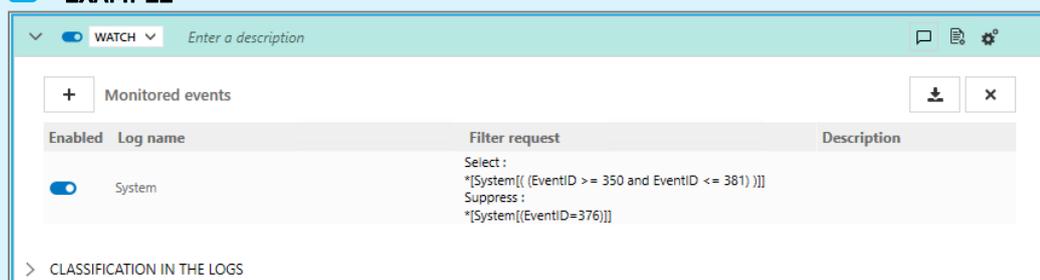
Here, the event IDs 350 to 381 in the System log will be forwarded, except for ID 376.

7. In the upper banner in the rule, you can:
   - Select the log settings that this rule will send. The severity of a log depends on its severity in Windows. Both severity levels are mapped as follows:

Windows event type	SES Evolution log
Audit	Information
Critical	Critical
Error	Error
Warning	Warning
Information	Information
Verbose	Diagnosis

   - Specify whether an action must be performed when a log is sent for this rule.
   - Enter a description to explain what this rule aims to achieve.
   - Enter a comment.

8. Add other event forwarding rules if necessary.

9. Click on **Save** at the top right of the window to save changes.

SES Evolution makes up for the Windows events that were generated when it was inactive, such as when the machine is restarting.

8.9.2 Importing OSSEC security rules

OSSEC is a host-based intrusion detection system, or a HIDS. It includes a monitoring and log analysis module. For more information, visit the OSSEC website.

SES Evolution is equipped with a similar analysis engine, which can monitor the following in real time:

- Log files from third-party applications,
- Windows events in event logs.

The aim of this type of monitoring is to extract information about SES Evolution agents in events and log lines, and to classify such information to identify abnormal or suspicious activity and generate alarms.

> ✏️ **EXAMPLE**
> You can monitor password-based authentication attempts on a FileZilla server from the same IP address, and raise alarms when there are multiple failures followed by a successful authentication.

> ℹ️ **NOTE**
> OSSEC analysis options will not be covered in detail in this document. Please refer to the relevant OSSEC documentation.

The Stormshield analysis engine and OSSEC differ in several ways:

- OSSEC collects logs on agents and analyzes them on the server while SES Evolution analyzes each agent. Events of the same nature occurring on separate agents therefore cannot be correlated.
- Unlike OSSEC, SES Evolution does not allow decoders and custom rules to be compiled. However, the rule *is_simple_http_request*, which OSSEC provides as an example but uses in standard configurations, is supported in SES Evolution.

For further information regarding all the OSSEC functions that SES Evolution supports, refer to **Supported OSSEC functions**

Configuring OSSEC rules

Configuring an OSSEC rule consists of indicating which log files and/or Windows events must be monitored and which decoder file and OSSEC rule to apply to them.

1. In an audit rule set, click on **External events > OSSEC rules**.
2. Click on **Add a rule (OSSEC)**.
3. If you want to monitor a log file from a third-party application, click on **+ Monitored file** and provide the following information:

Path
Enter the file path. You can use:

- Environment variables, only in the folder path up to the last \ of the path,
- File name specifications in *strftime* format only at the **end** of the path, after the last \ of the path.

> ✎ **EXAMPLE**
> If you enter the path *%PROGRAMFILES%\Filezilla Server\Logs\fzs-%Y-%m-%d.log*, SES Evolution will analyze any log line added to any file with a name in the form *fzs-YYYY-MM-DD.log*.

Encoding
Choose the type of encoding expected in the file. This depends on the application that generates logs. The supported encoding formats are:

- ANSI code pages, depending on the system locale,
- UTF8,
- UTF-16LE.

Description
Enter a description (optional). It will not impact the operation of the analysis in any way.

4. If you want to monitor a log or certain Windows events, click on **+ Monitored event** and provide the following information:
Log name
Enter the name of the Windows log, e.g., *System*, *Microsoft-Windows-Windows Defender/Operational*. To find out the name of a log, look up its properties in the Windows Event Viewer.

> **ⓘ NOTE**
> Logs that are not enabled in Windows can still be monitored. SES Evolution will automatically enable it. However, this operation may affect the performance of the host.

Filter request
If needed, enter a filter request to monitor only some events in the log. To obtain a request:

a. Open the Windows Event Viewer.

b. Right-click on the log of your choice > **Filter the current log**.

c. In the **Filter** tab, select your filtering options.

d. Copy the contents of the **XML** tab and paste it in the **Filter request** field in the OSSEC rule window.

Description
Enter a description if necessary. It will not impact the operation of the analysis in any way.

5. Click on **+ OSSEC decoder** and choose your *etc/decoder.xml* file. With an OSSEC decoder file, you can indicate which types of logs need to be analyzed and which values to extract. For more information, refer to OSSEC documentation
If you are importing several decoder files, ensure that they are in the right sequence, using the arrows on the left.

6. Click on **+ OSSEC rule sets** and choose your *etc-rules/*.xml files. Ensure that they are in the right sequence. The *rules_config.xml* file is mandatory and must be the first. It contains OSSEC rules 1 to 7 which must be the first rules declared.
You can also choose an OSSEC *.conf* file, in which case you must also specify the folder containing the rule files. Rules will be automatically imported in the same order.

7. Click on **Check the rule** to check the consistency of your OSSEC analysis configuration. The following aspects in particular will be checked:

• Validation of regular expressions found in the decoder files and rule files,

• Presence of decoders,

• Presence of rules 1 to 7,

• Validity of decoder files and rule files,

• Usage of OSSEC options that are not supported and therefore ignored.

The result of the verification shows errors, warnings and information messages:

• If errors are found, they will prevent the OSSEC configuration from being validated,

• Warnings will not prevent the configuration from being applied but may impact the evaluation of rules.

• Information messages indicate potential issues in the configuration and how they were resolved.

By default, the OSSEC analysis engine in SES Evolution retrieves Windows events generated when it is not enabled, e.g., when the machine is starting up. However, it does not retrieve log files.

Viewing logs generated by OSSEC

The external event logs that the SES Evolution analysis engine generates can be read like other SES Evolution logs in the administration console and on the agent. They are visible only to host administrators on the agent. For more information, refer to Viewing and managing agent logs in the administration console and Viewing logs in the agents' interface.

The logs contain all the fields collected during OSSEC decoding.

The severity of the log depends on the level of the OSSEC rule that specified the log:

Log level in the OSSEC rule	SES Evolution log severity
0	No log
1	Diagnosis
2	Information
3, 4, 5	Notice
6, 7, 8 and 9	Warning
10	Error
11, 12	Critical
13, 14	Alert
15	Emergency

✏️ **EXAMPLE**

The image below shows Filezilla logs extracted by the analysis engine and reported in the agent's interface. It detects password-based authentication attempts on a FileZilla server from the same IP address, and raises alarms when there are multiple failures followed by a successful authentication.
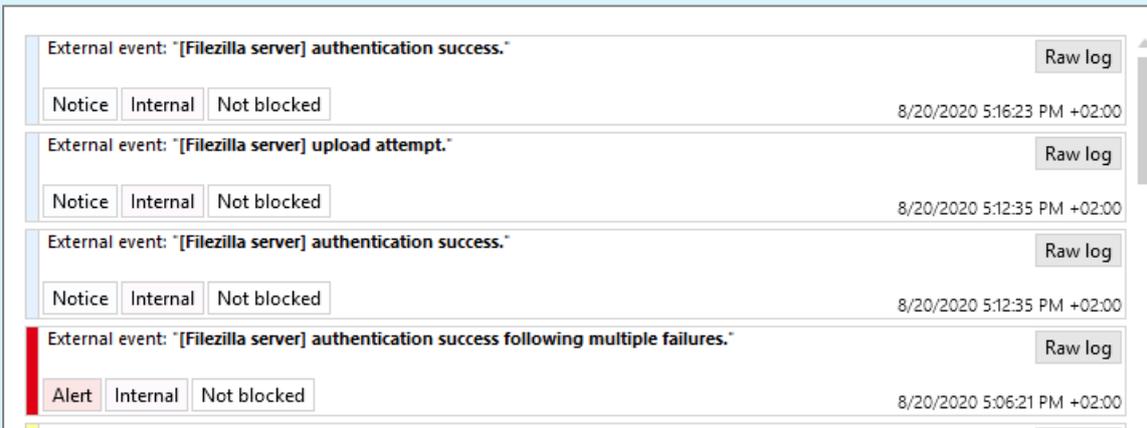
8.10 Testing security policies

We recommend that you test your security policies before deploying and implementing them on your pool.

By testing a policy, you will be able to measure the impact of usage restrictions that the policy places, and make adjustments accordingly to the protection rules before putting them into production.

The "Detection mode" and "passive rule" features make it possible to test policies on a pool without holding up the use of workstations, and testing takes place transparently for users. When these features are enabled, SES Evolution agents do not block operations, but instead, generate logs indicating the operations that would have been blocked by a rule.

It is helpful to test policies in the following cases:

- When you install SES Evolution for the first time on a pool of machines. In this case, testing a policy would allow you to know, for example, whether essential applications will be blocked, so that you can create suitable exceptions. Testing is transparent for users, so they can continue using their usual applications.

- When you expand the pool of machines to protect. You can create a new group of agents, for example, and test the application of the policy that is already implemented in the other groups. In this way, you can check whether the policy suits the new group before implementing it, and make adjustments where necessary.

- You need to add a new protection rule set to one of your security policies. Test the rule set first to check whether it blocks legitimate applications before putting it into production.

You can test security policies at various levels on a pool:

- the entire policy with Detection mode in the configuration of agent groups,
- an entire protection rule set with Detection mode in the configuration of a policy,
- a particular rule with the Passive rule mode in the configuration of the rule itself.

8.10.1 Testing an entire security policy assigned to an agent group

You can enable Detection mode in the configuration of an agent group. The configuration applies to all policies assigned to the group (main policy and conditional policies, if any).

By enabling Detection mode on an entire policy, this means that all rules from all protection rule sets will switch to Passive rule mode and the status of rules that filter threats will switch to "Detect only".

To test policies assigned to an agent group:

1. Select an agent group in the **Environment > Agents** menu.
2. In the **Policies** tab, enable the option **Switch policies to Detection mode**.
3. Deploy the environment.

This setting in the configuration of the agent group takes priority over the configuration on rule sets. As a result, if Detection mode is enabled for a group of agents, actions will be detected but not blocked, even if the rule set is in active mode in the policy.

For further information on the configuration of agent groups, refer to the section Creating agent groups.

8.10.2 Testing a protection rule set

Protection rule sets can be tested before they are put into production, regardless of whether they are private or shared. To test a rule set, you must enable the set's Detection mode in the relevant policy. In a shared rule set, even though Detection mode is enabled on the set to be tested, it will not be enabled on other policies that use the same rule set.

To enable Detection mode in a rule set:

1. Select the relevant policy in the **Security > Policies** menu.
2. Click on **Edit** in the upper banner.
3. In the row of the rule set to test, click on the arrow to the right of the shield icon ⛊ ⌄ .
4. Select the ⌗ icon to switch the rule set to Detection mode.

This setting applies to all agent groups that use this policy.

You can also fully disable a rule set with the 🔵 switch to the left of the rule. In this case, the rule set will not be deployed on the agent.

For more information on rule sets, refer to the section Understanding the difference between protection, exception and audit rule sets.

8.10.3 Testing rules

To find out what impact a security rule has without applying a block action, enable Passive rule mode in the options found in the upper banner of the rule.

Likewise, to find out what impact a rule has against threats without applying a block action, select "Detect only" as the status in the configuration of the rule.

For more details on these options and statuses, refer to the sections on the various rule types in Configuring threat protection and Defining access control rules.

8.11 Disabling security rules

All security rules can be disabled individually. Once a rule is disabled, the SES Evolution agent ignores it, as it is no longer part of the security policy.

Disable a rule if you want to stop using it temporarily without deleting it, or if you want to test the behavior of the agent without this rule.

Some types of rules are disabled by default when a new rule set is created. If such rules were enabled in several rule sets, they may overload and cancel the effect of the configuration due to the order of the sets. This is the case for threats, Wi-Fi connections, temporary web access and general devices.

To disable a rule:

1. Click on **Edit** in the upper banner.
2. In the top banner of the rule, deactivate the 🔵 switch. The rule is grayed out.

Disabling a rule is different from enabling **Passive rule** mode. For more information on passive rules, refer to the section Understanding the difference between protection, exception and audit rule sets.

8.12 Configuring log management

The agent generates logs whenever user actions are blocked or when the agent conducts an audit. Depending on their severity, these logs can be sent to three different destinations. The various settings of this process can be defined in the configuration of agent groups. For further information, refer to the section Sending logs generated by agents.

In addition, for every security rule that you create, you can specify:

- The severity of the logged events,
- The destinations of these logs.

> **ℹ NOTE**
> In any case, even if no destination has been configured for the logs in question, they can be found in the context details when an attack occurs. For further information on context analysis, refer to Understanding what makes up a context.

8.12.1 Recommendations

The severity of events logged by a rule can be adjusted in the following cases:

- If you have highly sensitive applications, raise the severity of their logs. *Emergency* and *Alert* logs take priority over other logs sent to agent handlers, and are sent more frequently (every 30 seconds by default, every hour for other log levels),
- If a security rule generates too many irrelevant logs, lower their severity.

8.12.2 Configuring logs in a security rule

1. Select your security policy in the **Security > Policies** menu of the administration console, then select your set of rules. The main page of the rule set appears.
2. Click on the tab of the rule that you want to modify.
3. If you are in read-only mode, click on **Edit** in the upper banner.
4. In the banner at the top of the rule, click on [icon]. The **Log settings** window appears.
5. In the **Log severity** field, assign the level to logs generated by this rule.
6. In the **Show on agent** field, choose whether logs from this rule can be seen on the agent:
   - **Inherit**: the overall behavior defined for the agent group applies. In the example above, logs can be seen on the agent because this is the case for logs of all levels from *Notice* upwards.
   - **Never**: logs can never be seen on the agent regardless of the overall behavior.
   - **Always**: logs can always be seen on the agent regardless of the overall behavior. Note that only *Alert* and *Emergency* level logs that have led to a block are visible in the agent interface for a non-administrator user of his machine.
7. In the **Show on console** field, choose whether logs from this rule can be seen on the administration console.
8. In the **Send to Syslog** field, choose whether to send logs from this rule to the Syslog server if one has been configured. For further information, refer to the section Creating groups of agent handlers.
9. Click on **OK**.
10. Save the changes made to the rule.

8.13 Configuring actions triggered by rules

When a protection rule blocks an operation performed on an SES Evolution agent, it will be logged, and you can determine the severity and destination of the log.

If you want it to, when this log is generated, it can trigger other actions on the agents in question. Various types of actions are possible:

- Show a notification on the agent. This notification will appear at the bottom right of the screen, indicating that a prohibited action was blocked by a protection rule.
- Run custom scripts.
- Run a Yara or IoC scan. For more information, refer to Running Yara scans and Searching for indicators of compromise
  This action is available only for rules in which processes or files have been logged. It does not apply, for example, to Devices and ARP Spoofing rules.

> ✏️ **EXAMPLE**
> This feature may be useful in triggering an antivirus analysis the moment the incident is logged, or it can move a dangerous file to a specific folder. Launching a Yara or IoC scan makes it possible to identify malware, for example.

1. Select your security policy in the **Security > Policies** menu, then select the set of rules. The main page of the rule set appears.
2. Click on the tab of the rule that you want to modify.
3. If you are in read-only mode, click on **Edit** in the upper banner.
4. In the banner at the top of the rule, click on ⚙️. The window **Action when logs generated** appears.
5. Enable a notification on the agent, if you wish to, for every time this rule triggers a log. This feature is available only for rules in Protection mode.
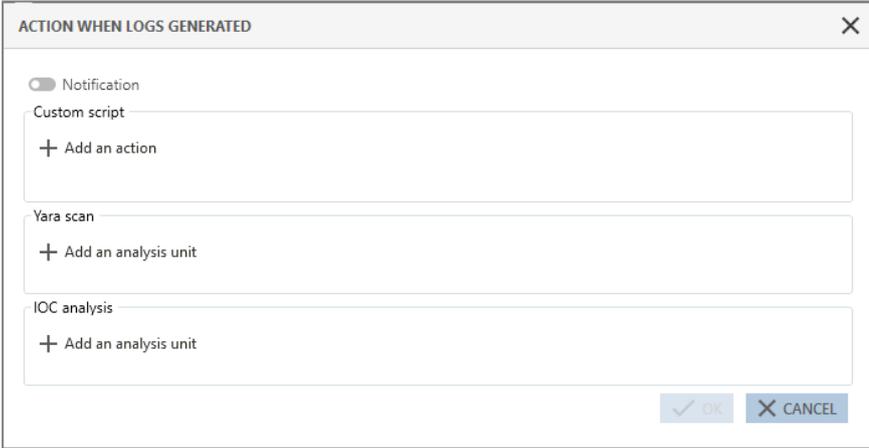
6. If you wish to run a script whenever this rule generates a log, click on **Add an action**.
   a. Enter a name for the action in the **Run custom script** window.
      a. To the right of the **Script** field, click on + to add the script to run.
      b. In the **Arguments** field, specify the arguments to add when the script is run.
      c. In the **Run in** list, choose **Local service** because this is an account with restricted privileges. Do not choose **Interactive session** or **System** accounts unless absolutely necessary.

      Do note that scripts cannot be run during interactive sessions on a server with several remotely connected users.

All scripts that were declared in SES Evolution appear in the **Script** list. Select an existing script and click on ⬚ to view it or ⬚ to import a new version of the script.

7. If you wish to run a Yara or IoC scan whenever this rule generates a log, click on **Add an analysis unit** in the relevant section.

   a. Click on one or several analysis units to select them, then close this window. For an IoC scan, only text or file name indicators can be used here.

   b. In the **Action when logs generated** window, click on Log settings to determine the severity and destination of the logs that the Yara or IoC rules generated.

   c. If necessary, select **Shut down the processes detected** to remove the dangerous processes, identified during the Yara or IoC scan, from the agent.
   If the rule is part of an audit rule set, or if the rule is in passive mode, the processes will not be shut down even when this setting is enabled.

   d. If the above option is enabled, select **Quarantine** if you also want to quarantine detected files. For more information, see the section Managing file quarantine.

8. Click **Validate**.

For IoC scans activated when a log is generated, they can analyze only the element that triggered the rule. Scheduled or on-demand IoC scans are more precise as their settings make it possible to include or exclude the folder to scan. For more information, refer to Scheduling IoC scans and Running IoC scans on demand.

> ✏️ **EXAMPLE**
> In a File rule that protects files located in *C:\temp* from being deleted, you configured an IoC scan when logs are generated, to search for "*suspect_text*" in files. If the user attempts to delete a file in *C:\temp* with PowerShell, the rule will apply and the IoC scan will be triggered, but only on the *powershell.exe* file and on the memory of the *powershell.exe* process.

8.14 Assigning a security policy to agents

The Stormshield Default Policy is applied by default to agent groups, but customized security policies can also be assigned to agent groups.

1. Select the **Environment > Agents** menu.

2. Select an agent group from the left panel.

3. In the upper banner, click on **Edit**.

4. Go to the **Policies** tab in an agent group.

5. Choose the security policy that you want to apply to all agents in the group from the **Policy** drop-down list.
   For further information, refer to the section Creating agent groups.

6. In the upper banner, click on **Save**.

7. To deploy the policy on all agents in the group so that they apply this policy, go to the **Security > Deployment** menu and click on **Deploy**.

8.15 Importing and exporting policies and rule sets

Full policies or only rule sets can be exported to a *.cab* file, itself containing *.json* files. This file can then be reimported. This makes it possible to:

- Transfer a security policy from a pre-production environment to a production environment,
- Transfer a policy or rule set to SES Evolution's technical support to make it easier to debug issues.

When you export a policy or rule set, you export the version selected in the right side of the panel, represented by the 📍 icon. When you import a custom policy or rule set that already exists, its version number will be incremented each time.

For further information on how to manage versions of policies and rule sets, refer to the section **Managing versions of a policy or a rule set**.

8.15.1 Exporting all security policies in the list

1. In the **Security > Policy** menu, click the **Export-Export all** button of desired policy type.
2. Select the folder to which you want to export the file.
   The latest version of each policy will be exported in the form of an individual file named *policy_name.cab*.

8.15.2 Exporting one or several security policies

1. Select the policy(ies) to export in **Security > Policies**.
2. By default, the latest version of a policy will be exported. If you wish to export another version, select it in the right column of the policy's general panel.
3. Click on **Export-Export the selection** and choose the folder to which you want to export the file.
   Each policy will be exported in the form of an individual file named *policy_name.cab*.

8.15.3 Importing one or several security policies

- To import a single policy or several policies in one go, select **Security > Policies** and click the **Import** button of the desired policy type.
  Importing an existing policy will automatically create a new version of this policy, unless it is a built-in policy.

8.15.4 Exporting rule sets

1. Double-click on the policy in **Security > Policies**, then select a set of rules.
2. By default, the latest version of rule set will be exported. If you wish to export another version, select it in the right column.
3. Click on **Export** and choose the name of the file and the folder to which you want to export the file.

8.15.5 Exporting a selection of shared rule sets

1. In the **Security > Policy** menu, click the **Shared rule set** button of the desired policy type.
2. Select the rule sets that you want to export, and/or filter the list of rule sets by using the search field at the top right side.
3. Click on **Export-Export all** to export all the rule sets shown in the list, or **Exporter-Export the selection** if you have selected only a few rule sets.

4. Select the folder to which you want to export the file.
   Each rule set will be exported in the form of an individual file named *set_name.cab*.

You cannot export several private rule sets at the same time.

8.15.6 Importing rule sets

1. Double-click on the policy of your choice in **Security > Policies**.
2. In the general panel of the policy, click on **Import** and choose the *.cab* file of the rule set(s) that you want to import.
   Importing an existing rule set will automatically create a new version of this set.
   If it is a built-in rule set, it will create a new version only if it does not already exist.

9. Deploying the SES Evolution environment

To apply the configuration of agent groups, security policies and new software versions of the agent to your pool of agents, the environment must be deployed.

Doing so will generate for each agent group the information to send to agents. Configuration and policy packages are generated and stored in databases. Agents log in regularly to their agent handlers to update their statuses. The agent handler then detects updates to apply to agents when they become available.

By default, agents connect to their handlers every 60 seconds by default. This means that it takes less than one minute for a new deployment to be applied. This duration can be modified in the configuration of agent groups using the **Agent status update** setting. For further information, refer to the section Monitoring agents in real time.

The environment must be deployed again every time you make changes in the administration console to the following items and you wish to apply them to the pool:

- Policies and rule sets,
- Agent group configuration,
- Agent handler configuration,
- Configuration of a USB key's trustworthiness.

Ensure that you hold the **Environment-Deploy** privilege to perform this action.

To deploy the environment on agents in the pool:

1. Change one or several of the configuration items mentioned above.
   An orange dot appears to the right of the **Security > Deployment** menu and the **Deployment** icon in the upper banner of the console turns orange . This means that new elements must be deployed in order to be operational.
2. Click on the **Deployment** icon, or select the **Security** > **Deployment** menu, and click on **Deploy**.
   The orange dot disappears and the **Deploy** button turns gray until the next time the configuration is edited.

> **ⓘ NOTE**
> The more cores the machine hosting the backend component has, the shorter the deployment time, especially when you have configured a large number of agent groups. The minimum recommendation is two cores. For more information on system requirements, see the *Backend* section of the *Installation guide*.

As for the versions of rule sets deployed with a policy, you can select **Always use latest version** in the policy's general panel. In this case, after the policy is deployed, the version number of the deployed rule sets appears in the drop-down list. When you click on **Edit** in this panel, the **Always use latest version** setting is retained for each set. For further information, refer to the section Managing versions of a policy or a rule set.

If the environment cannot be deployed, the interface shows a message providing the reason or the actions that must be performed before it can be deployed.

The environment can be deployed from the console only on agents connected to agent handlers. To apply configuration or software updates to agents that are not connected to agent handlers, refer to the section Updating agents.

10. Managing devices

SES Evolution allows you to control access to all types of devices that can be connected to users' workstations, based on their type, trustworthiness, content, etc.

The following table sets out the list of protections that apply to each device type, and the security rules that allow them to be configured.

Device	I'd like to:	I need to use:
USB	Filter the use of some types of USB devices based on their characteristics, e.g., class, vendor, serial number, etc. **Example:** Allow only wireless USB mice issued by the organization, or prohibit the connection of any USB key.	Access control rules for USB devices in **Security > Policies**, Peripherals > USB rules.
	Block access to any unknown device that has never been monitored by a decontamination station.	• The configuration of agent groups in the **Environment > Agents > Policies > Trusted devices** menu. • The control rules for storing data on USB devices in the **Security > Policies** menu, Devices > USB storage rules. • The control panel for trusted USB devices in the Security > Devices menu.
	Filter access to data on a USB mass storage device. **Example**: Allow access only to office files.	• The access control rules for files in the **Security > Policies** menu, ACL resources > File rules. Select the **Removable** option in the identifier's volume type. - or - • The control rules for storing data on USB devices in the **Security > Policies** menu, Devices > USB storage rules.
	Filter the execution of an application from a removable mass storage device. **Example**: Allow only a specific software program from the IT department to run.	• The application identifiers in the **Security > Policies** menu, Identifiers > Application identifiers rules. Select the **Removable** option in the volume type. - or - • The control rules for storing data on USB devices in the **Security > Policies** menu, Devices > USB storage rules.
Bluetooth	Filter the use of some types of Bluetooth devices based on their class. **Example:** Allow only Bluetooth headsets issued by the organization.	The access control rules for Bluetooth devices in the **Security > Policies** menu, Devices > Bluetooth rules.
CD/DVD	Filter the use of CDs and DVDs.	The access control rules for general devices in the **Security > Policies** menu, Devices > General rules.
Floppy disk	Filter the use of floppy disks.	
Serial port	Filter the use of devices on serial ports.	

10.1 Controlling access to devices

SES Evolution allows you to control access to all types of devices that can be connected to users' workstations, based on their type, trustworthiness, content, etc.

10.1.1 Controlling access to general devices

This protection type allows you to control how floppy disk drives, CD/DVD drives and serial ports are used on physical or virtual user workstations. Floppy disk drives and serial ports are found mostly in industrial environments.

For every type of device, you have the option of allowing, blocking (in a protection rule set) or simply monitoring its use (in an audit rule set).

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **Devices > General** tab. Access to all devices is allowed by default and rules are disabled. Enable them by clicking on [toggle] on the left if you want to block access (Protection mode) or monitor access (Audit mode). Ensure that your rule sets are in the right order if these rules are enabled in several rule sets, as they may overload and cancel the effect of the general device access configuration in the rule sets that follow.

4. For every device type, select the action to apply whenever the device is used or plugged in. For the **CD/DVD drive** type, the **Read only** option allows access to the device but blocks burning.
   If you select the **Block** or **Audit** action, a log will be generated only when the device is first used.

5. In the banner at the top of the rule:

   • Select the log settings that this rule will send.

   • Specify whether an action must be performed when a log is sent for this rule.

Floppy disks or CD/DVDs inserted into external USB drives, and serial ports linked by a USB cable are considered both USB devices and floppy disk or CD/DVD drives, or internal serial ports. They can therefore be blocked either from the **General** tab or the **USB** tab.

10.1.2 Controlling access to Bluetooth devices

This protection type allows you to control how Bluetooth devices are used on user workstations.

SES Evolution makes it possible to monitor when Bluetooth devices are connected and disconnected, by generating logs if Audit mode is enabled in an audit rule set. Access to Bluetooth devices can also be blocked in a protection rule set.

Security rules can be configured to filter Bluetooth devices based on their class. To understand Bluetooth classes, refer to the IEEE standard on Bluetooth.

> **ℹ NOTE**
> If a multifunction Bluetooth device is blocked by a rule, all of its functions will be blocked. For example, if a rule blocks the use of the microphone class, headsets will also be blocked.

To create rules for Bluetooth devices:

1. Select the **Security > Policies** menu and click on your policy.

2. Select a rule set.

3. Click on the **Devices > Bluetooth** tab.

4. If you are in read-only mode, click on **Edit** in the upper banner.

5. Click on **Add > Rule (Bluetooth devices)**. A new line is displayed.

6. On the left side of the rule, click on ☐☐☐ to add Bluetooth device identifiers.

7. Enter a name for each identifier.

8. Select the device's service class and major class.

9. Click on **OK**.

10. In the **Access** field, select **Allow** or **Block** if you are in a protection rule set, or **Allow** or **Audit** if you are in an audit rule set. **Skip rule set** allows you to ignore all the rules contained in this rule set and evaluate the next rule set.

11. In the upper banner in the rule, you can:

    • Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
    Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

    • Select the log settings that this rule will send.

    • Specify whether an action must be performed when a log is sent for this rule.

    • Enter a comment.

    • Enter a description to explain what this rule aims to achieve.

12. The row number of each rule appears on its left. Rearrange the sequence of your rules if you need to, by clicking on the arrows above and below the row number.

13. Click on **Save** at the top right of the window to save changes.

If you only want to monitor the use of Bluetooth devices in the pool:

1. Create a Bluetooth device rule in an audit rule set.

2. Create an identifier that includes all Bluetooth device classes.

3. Select **Audit** as the action in the **Access** field.

4. Analyze logs that are generated every time a device is connected and disconnected.

10.1.3 Controlling access to USB devices

This protection type allows you to control how USB devices are used on user workstations.

Rules may apply to USB device classes (printer, video, audio, storage, etc.) and/or vendors, models or device serial numbers.

For every USB device category, you can:

• Authorize their use,

• Block its use,

• Display a message for the user to confirm whether or not to use the device when it is connected. The confirmation request is shown to the user who is currently connected locally on the workstation. In other situations when the device is plugged in (e.g., remote connection, local session locked, startup or session signed out), the confirmation message does not appear and the device will always be blocked.

• Monitor the use of USB devices in a set of audit rules.

> ✎ **EXAMPLE 1**
> SES Evolution also allows the detection of *Rubber Ducky* USB keys. Such keys act as keyboards, run malicious scripts and save data on micro SD cards. If you create a rule that asks for user confirmation every time an HID or human-interface device (e.g., keyboards or mice) is plugged in, a message will indicate that a keyboard has just been plugged in. The user can then deny access to this malicious device that appears to be a USB key.

> ✎ **EXAMPLE 2**
> You can choose to allow only headsets, speakers and mobile phones provided by your company's IT department.

> ❗ **CAUTION**
> When keyboards or mice (HID peripherals) are plugged in before the workstation has started, they are automatically allowed so that they do not render the workstation unusable.
> However, on Microsoft Windows 10 and 11 operating systems, if **Turn on fast startup** is selected, it will switch to Hibernate mode when the PC is shut down. This means that if you have set a rule on an HID containing a block action or request for user confirmation, the device will be blocked once the PC comes out of hibernation. You must then fully restart the PC by using the standard menus, or by holding down the power button.

If you choose to apply a whitelist, you must create rules to allow the use of certain devices in your pool. The last rule must block all other devices. We recommend that you choose the Passive rule mode for the last rule to avoid blocking devices that allow workstations to run properly. Doing so will allow you to test the rules you want to apply to USB devices in a production environment, and refine them later after checking the logs.

To create rules for USB devices:

1. Select the **Security > Policies** menu and click on your policy.
2. Select a rule set.
3. Click on the **Devices > USB** tab.
4. If you are in read-only mode, click on **Edit** in the upper banner.
5. Click on **Add > Rule (USB device)**. A new line appears.
6. In the left side of the rule, click on [=∥+] to indicate one or several device identifiers to which the rule applies. Click on **Select an existing device** to view the list previously defined in the **Security** > **Device** menu, otherwise fill in the fields manually. Depending on whether you want to filter a specific device or a device category, fill in some or all of these properties:
   - Enter a name for this device,
   - Select the USB class of the device from the drop-down list. Click on [✎] to enter a value manually if necessary.
   - Enter the USB sub-class consisting of two hexadecimal characters.
   - Enter the first few letters of the vendor name to show the list and select the desired vendor. You can also enter the four standardized hexadecimal characters corresponding to the vendor.
   - Select the product from the list of this vendor's products or enter the four hexadecimal characters.
   - Enter the serial number of the product.

7. In the **Access** field, select **Allow**, **Block** or **Request** if you are in a protection rule set, or **Allow** or **Audit** if you are in an audit rule set. **Skip rule set** allows you to ignore all the rules contained in this rule set and evaluate the next rule set.

8. In the upper banner in the rule, you can:
   - Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
     Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.
   - Select the log settings that this rule will send.
   - Specify whether an action must be performed when a log is sent for this rule.
   - Enter a comment.
   - Enter a description to explain what this rule aims to achieve.

9. The row number of each rule appears on its left. Rearrange the sequence of your rules if you need to, by clicking on the arrows above and below the row number.

10. Click on **Save** at the top right of the window to save changes.

To find out the vendor or product IDs, or the serial numbers of devices, look up the Windows device manager when the device in question is plugged in or use the dedicated utilities.

Refer to the international USB standard to find out the identifiers of USB device sub-classes.

10.1.4 Controlling storage on USB devices

This protection makes it possible to control access to files stored on USB storage devices (external hard disks, USB keys, etc.).

Rules may cover devices filtered by vendor ID or product ID, or devices known to SES Evolution with a trust level.

If overall access to USB devices is blocked, files on USB mass storage devices cannot be accessed even when a rule specifically applying to these devices allows it. To monitor overall access, refer to the section Controlling access to USB devices.

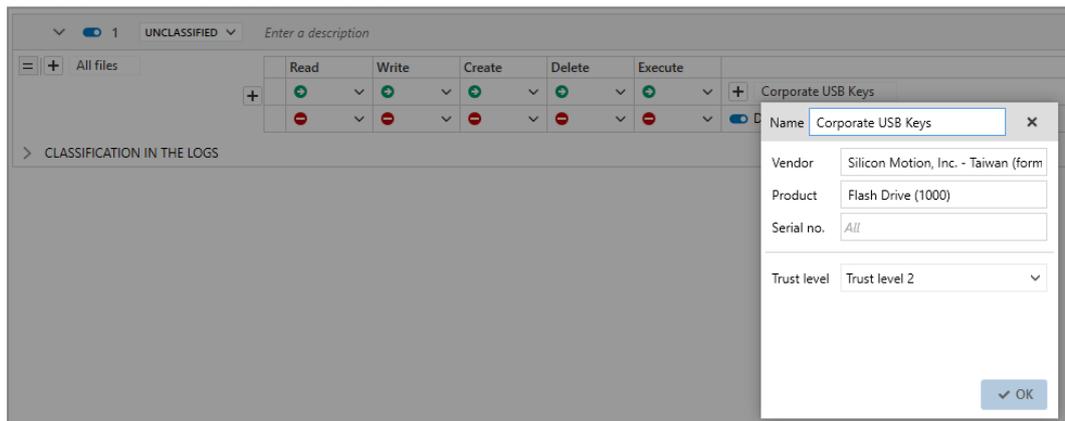For more information on trust levels, refer to the section Managing USB storage devices.

The left side of a rule covers files that may be found on USB devices, while the right side covers the devices themselves.

To create rules that regulate access to files on USB storage devices:

1. Select the **Security > Policies** menu and click on your policy.
2. Select a rule set.
3. Click on the **Devices > USB storage** tab.
4. If you are in read-only mode, click on **Edit** in the upper banner.
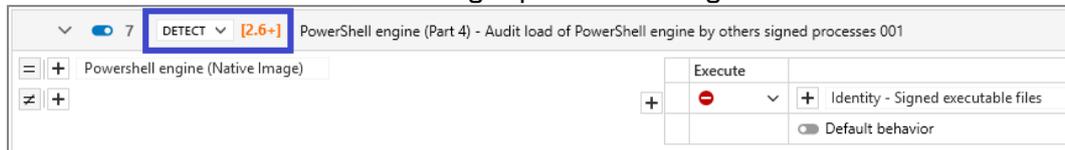5. Click on **Add > Rule (USB storage device)**. A new line appears.

6. In the left side of the rule, click on [= +] to add file identifiers. Files can be identified by a path or an alternate data stream. Generic characters are allowed in this field.

7. Click on **Apply** to add the ID.

8. For each type of operation, select the default behavior that applies to the devices when files match the rule: allow or block (protection rule).
   You can also:

   - **Skip behavior** to ignore the subrule if the behavior is detected and move on to the next behavior.

   - **Skip rule set** to ignore all the rules contained in this rule set and evaluate the next rule set.

9. To exclude specific devices from the default behavior, click on **+ Add specific behavior**:

   a. Add one or several device IDs. Devices can either be identified by their vendor or product IDs, or the trust level that SES Evolution assigned to the device. Click on **Select an existing device** to view the list previously defined in the **Security** > **Device** menu, otherwise fill in the fields manually.

      - To find out the vendor or product IDs, or the serial numbers of devices, look up the Windows device manager when the device in question is plugged in or use the dedicated utilities.

      - For more information on trust levels, refer to the section Managing USB storage devices.

   b. Select the behavior for these IDs.

10. In the upper banner in the rule, you can:

- If necessary, rearrange the order of the rules by clicking on ⊞ when the cursor hovers the rule. Each rule displays its line number in the banner.

- Disable rule. For more information, refer to the section Disabling security rules.

- Indicate the intent of the rule, according to predefined categories:

- ○ Unclassified: unclassified rule.

- ○ Nominal: non-blocking rule conforming to nominal application behavior.

- ○ Protect: blocking rule with a high log severity level.

- ○ Protect silent: blocking rule with a severity level below the log thresholds displayed by default on the agent and console. Protects access to resources deemed sensitive, even if carried out by programs with no malicious intent. As there may be many such programs, a rule with too high a log severity could trigger massive log generation.

- ○ Detect: non-blocking audit rule or passive rule.

- ○ Context: rule used to build an attack graph.

- ○ Syslog: rule triggering logs sent exclusively to a Syslog server.

- ○ Watch: rule for monitoring behavior in order to fine-tune the security policy or gain a better understanding of technical events occurring in the pool.

Selecting one of these categories has no influence on rule configuration. They simply enable the administrator to classify their security rules according to their purpose, and sort them using the dedicated **Rule intent** filter. The rule intent is also displayed in the log details.

- Enter a description to explain what this rule aims to achieve.

- Make the rule passive. Passive rules behave like standard rules but do not actually block any actions. The agent only generates logs that indicate which actions security rules would have blocked.
  Use this mode to test new restriction rules, determine their impact, and make the necessary adjustments before disabling **Passive rule** mode. For further information on testing rules and policies, refer to Testing security policies.

- Indicate whether the rule must generate a context when it is applied. By default, if a rule generates *Emergency* or *Alert* logs, it will generate a context, but you can disable this feature. In case of mass generation of similar logs, the context is not generated. For more information on mass log generation, refer to the section Monitoring SES Evolution agent activity.

- Adding a comment.

- Select the log settings that this rule will send.

- Specify whether an action must be performed when a log is sent for this rule. You can request that a script be run and/or that a Yara or IoC scan be triggered. You can also request that a notification be displayed on the agent, provided that it is associated with a an *Alert* or *Emergency* level blocking log.

- Deleting the rule.

11. Expand the **Classification in logs** part to indicate the intent of the suspected attack when the rule applies, along with the tags for associating the rule with the MITRE repository. This information is then visible in the logs generated by the rule. For more information, see Classifying attacks according to the MITRE repository.

12. Click on **Save** at the top right of the window to save changes.

## 10.1.5 Controlling application execution from removable devices

SES Evolution makes it possible to control the execution of applications found on USB storage media. Two methods are available depending on the use case:

- Use case 1: I want to request confirmation from users when they attempt to run an application on a USB storage device.
- Use case 2: I want to allow the execution of applications only from a certain type of USB key that the company provides to employees. These keys are identified by their vendor IDs and product IDs and/or trust level.

Both of these use cases can also be combined.

Requesting user confirmation

1. Create an application identifier that indicates:
   - The applications for which you want to request confirmation. Type the **Path** *.exe* for example to indicate that all applications are concerned.
   - The type of volume in question. Enable only **Removable** in this case.
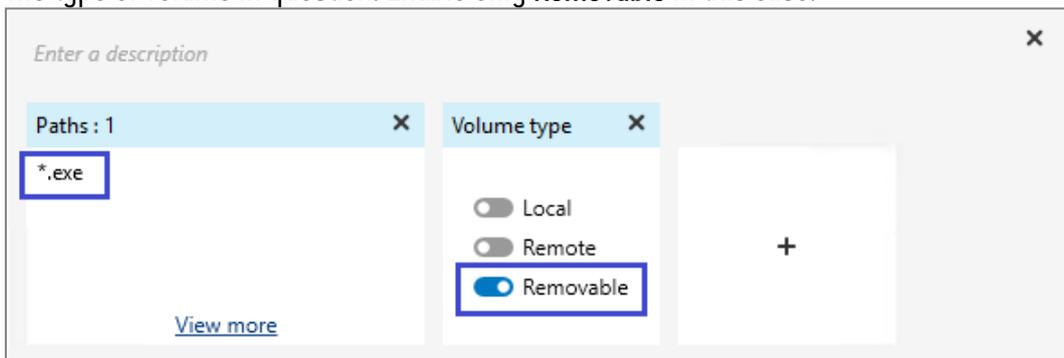
For more information, refer to the section Creating application identifiers.

2. Create a process creation rule that indicates:
   - The application identifier created above,
   - That users must confirm whenever they execute applications from a removable device. Select **Request** as the default behavior.
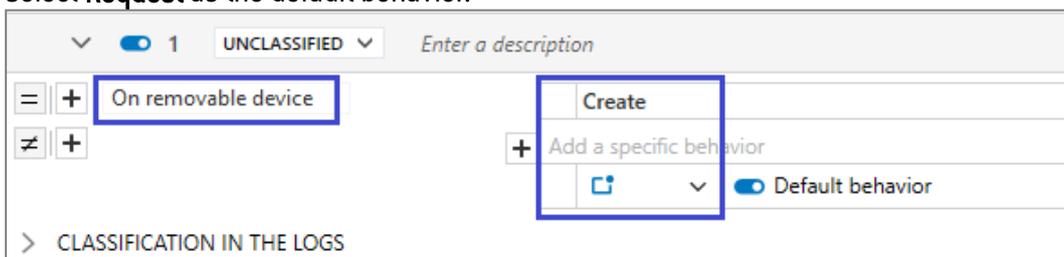
For further information, refer to the section Controlling process creation.

Once this rule is created, the user will be able to run applications from removable devices only after confirming that the action is deliberate. The request for confirmation and the user's response will be logged in the agent.

Allowing application execution only for certain key types

Create a USB storage rule that indicates:

- The application(s) that you want to prohibit if they are found on a USB storage device. In the section on the left, type the **Path** *.exe for example to indicate that all applications are concerned.
- The desired default behavior. Choose **Block** from the **Execution** drop-down list to block the execution of applications.
- The type of keys on which applications are allowed to run. In the right side of the rule, enter the hardware information of this type of key and/or the desired trust level.

For further information, refer to the section Controlling storage on USB devices.

Once this rule is created, applications will be prohibited from running on USB storage devices except for trust level 2 devices.

10.2 Managing USB storage devices

With SES Evolution, USB storage devices such as external hard disks and USB keys can be monitored. In this section, the term *USB device* is used to refer to such devices.

Whenever a USB storage device is connected to an SES Evolution agent, a log is generated and appears in the **Devices** panel in the administration console if detection options have been enabled in the agent's group. In this panel, you will be able to view all USB devices that were plugged into your appliance pool and find out their level of trust. You can also modify the level of trust of various devices and manually pre-declare devices.

SES Evolution does not support USB devices that contain several partitions. Even though such devices appear in the list of devices, the information regarding them is inaccurate.

Some operations on devices can be automated for an agent group. For further information, refer to the section Detecting and configuring the trust level on devices

Depending on whether you want to make changes or only view the **Devices** panel, you must have the **Removable devices-Modify** or **Removable devices-Display** privilege.

10.2.1 Viewing USB devices

To ensure that plugged in USB devices on agents appear in the console, detection options must be enabled beforehand in agent groups. For further information, refer to the section Detecting and configuring the trust level on devices

1. Select the **Security > Devices** menu. You will see the list of all USB devices that have ever been plugged in when SES Evolution agents are used.

2. Refer to the specific information about devices. In addition to the name, size and hardware information, the following details are provided:
   - **Status**: new or modified,
   - **Workstation**: last workstation on which the device was plugged in,
   - **Session**: user session opened the last time the device was plugged in,
   - Trust level,
   - **Last seen**: date of the last time the device was plugged in on an SES Evolution agent,
   - **Unique ID**
   - **First seen**: date of the first time the device was plugged in on an SES Evolution agent.
   - **Description**: custom comments that you added when you modified the device in SES Evolution,

3. Not all columns are displayed by default. To show additional columns, right-click on the row of column headers and select the ones you want to display.

4. In the **Filters** area, select the USB devices that you wish to show in the list by filtering them according to their **current** and/or **desired** trust level. Click on **Reset filters** at the top right side to display all USB devices again.

10.2.2 Adding a description to a USB device

1. Select the **Security > Devices** menu. You will see the list of all USB devices that have ever been plugged in when SES Evolution agents are used.

2. Select one or several devices and click on **Change selection**.

3. Click on [pencil icon] in **Description** to add a comment.

4. In the **Trust level** area, select **Keep trust level** as the action.

5. Click **OK**. The added comment will appear in the **Description** column of the USB device.

10.2.3 Changing the trust level of a USB device

There are three trust levels for USB devices in SES Evolution:

- **Level 0** - [icon]: For the SES Evolution agent, the device is neither enrolled, nor trusted. The device is plugged into the SES Evolution agent but the backoffice has not yet assigned an unique ID to it.

- **Level 1** - [icon]: For the SES Evolution agent, the device is enrolled, but not trusted. The device is known and the backoffice has assigned a unique ID to it. Either its content has not yet been verified or it has changed since the last verification (when changes are made to a host outside the SES Evolution pool, for example). The device must be analyzed by an air-gapped workstation to switch to level 2.

- **Level 2** - [icon]: For the SES Evolution agent, the device is enrolled and trusted. The device is known to the backoffice with a unique ID and its content is considered trusted. This level indicates that the device has been analyzed by an antivirus on an air-gapped SES Evolution workstation and that it does not contain any malicious files. This trust level will be maintained as long as the device's content is changed within the SES Evolution pool.

The trust level of a device is recognized throughout your SES Evolution pool, and does not depend on agent groups.

Once the trust levels are assigned, use them to filter the USB devices allowed in your pool. For example, you can protect your pool by creating a rule that allows only level 2 USB devices. For further information, refer to the section Controlling storage on USB devices.

For security reasons, the trust level of a USB device cannot be changed in the following cases:

- If the user session on the agent is locked or signed out,
- If the agent is remotely controlled through a remote desktop connection,
- If the device was already connected when the agent started running.

To change its trust level, the device must be inserted after the user session is opened on the physical workstation.

Granting trust level 1 to USB devices

1. Select the **Security > Devices** menu. You will see the list of all USB devices that have ever been plugged in when SES Evolution agents are used.
2. Select one or several devices and click on **Change selection**.
3. In the **Trust level** area, select **Raise the trust level of level 0 devices** as the action.
4. Click on **OK**.

   The change in trust level is shown in the corresponding column in the **Devices** panel. The icon means that the level 0 device will switch to level 1 the next time it is connected to an SES Evolution agent.
5. To apply this change to agents, select the **Security > Deployment** menu and click on **Deploy**.
6. Connect the modified device to an SES Evolution agent (or disconnect and reconnect it if it had stayed connected). It will appear in the panel of devices with its new trust level 1 .

Level 1 can also be automatically granted to any device that is connected to an SES Evolution agent if the option **Allow device identification** was enabled in the configuration of the agent group. For further information, refer to the section Detecting and configuring the trust level on devices.

Granting trust level 2 to USB devices

Trust level 2 can only be granted after the USB device has been connected to a decontamination station. A decontamination station is a dedicated SES Evolution agent on which USB devices in the pool are analyzed and granted the highest trust level if they are considered trustworthy. In general, it is equipped with one or several antiviruses that are more powerful than the other agents in the pool, and a specific SES Evolution security policy.

1. Configure your SES Evolution agent as a decontamination station:
   - Add it to an agent group in which it will be the only agent.
   - Configure the agent group by enabling the **Trust empty devices** and **Automatically scan devices** options.
   - Deploy the policy on the agent from the **Security > Deployment** menu.

2. Plug the USB device into the decontamination station.
If it is considered trustworthy, it will appear directly in the **Devices** panel with the highest trust level. It will lose this trust level as soon as its contents are modified outside the SES Evolution pool. Plug it into the decontamination workstation again to restore the highest trust level.

Untrusting USB devices

Untrusting a USB device means that its trust level will be brought down to 0.

1. Select the **Security > Devices** menu. You will see the list of all USB devices that have ever been plugged in when SES Evolution agents are used.

2. Select one or several devices and click on **Change selection**.

3. In the **Trust level** area, select **Untrust level 1 or 2 devices** as the action.

4. Click on **OK**.

The change in trust level is shown in the corresponding column in the **Devices** panel. The icon means that the level 1 device will switch to level 0 the next time it is connected to an SES Evolution agent.

5. To apply this change to agents, select the **Security > Deployment** menu and click on **Deploy**.

6. Connect the modified device to an SES Evolution agent (or disconnect and reconnect it if it had stayed connected). It will appear in the panel of devices with its new trust level 0 .

10.2.4 Pre-declaring USB devices

When USB devices are pre-declared, they will be easier to identify. When they are connected to an SES Evolution agent, they will be enrolled - the agent will recognize these devices and automatically grant them the predefined trust level. For example, you can pre-declare devices distributed to coworkers in order to automatically assign trust level 1 to them the first time they connect to an SES Evolution agent.

For security reasons, SES Evolution does not enroll connected USB devices in the following cases:

- If the user session on the agent is locked or signed out,
- If the agent is remotely controlled through a remote desktop connection,
- If the device was already connected when the agent started running.

To enable the enrollment of a device, it must be inserted after the user session is opened on the physical workstation.

To pre-declare USB devices:

1. Select the **Security > Devices** menu.

2. Click on **Add**.

3. Specify the **Vendor** and **Product** IDs of the device.

4. Enter its **Serial number** and a **Description** if you wish to (optional).

To find out these identifiers or the serial numbers of devices, look up the Windows device manager when the device in question is plugged in or use the dedicated utilities.

5. Choose the **Trust level** that will automatically be assigned to this device when it is connected to an SES Evolution agent: Trust level 0 or 1.

6. Click on **OK**.
   A line corresponding to this new device appears in the **Devices** panel. You will see only the information that you have specified.

7. To send information about pre-declared devices to agents, select the **Security > Deployment** menu and click on **Deploy**.
   When the device is connected to an SES Evolution agent, it will be identified and information found in the **Devices** panel will be filled in.

10.2.5 Removing USB devices

1. Select the **Security > Devices** menu.

2. Right-click on the USB device you want to remove, and select **Delete**.
   The device will no longer appear in the list.

   The next time it is connected to an SES Evolution agent, the device will appear once again in the **Devices** panel with the same level of trust that it had when it was deleted.

10.2.6 Importing and exporting a list of USB devices

You can import or export a list of USB devices in a CSV file in the **Security > Devices** panel.

Importing a list of USB devices

1. Select the **Security > Devices** menu.

2. Click on **Import** and select the file to import.
   The file must be in CSV format with one device per line. The syntax is as follows:
   `Product ID, Serial number, Vendor ID, Trust level, Description`

   Product ID, vendor ID (in four-character hexadecimal format) and confidence level (0, 1 or 2) are required. Note that level 2 is automatically converted to level 1 when imported into the administration console. For more information on trust levels, see Changing the trust level of a USB device.

> 🖉 **EXAMPLE**
> The line `5834,,0A5C,2,Stormshield key` will import a key with the following properties:
>
> | | |
> |---|---|
> | **Product ID** | 5834 |
> | **Serial number** | not entered |
> | **Vendor ID** | 0A5C |
> | **Trust level** | 1 |
> | **Description** | Stormshield key |

Exporting a list of USB devices

1. Select the **Security > Devices** menu.
2. To export all the devices in the list, click on **Export**.
   - or -
   To export only some devices, select them from the list, then click on the arrow of the **Export > Export the selection** button.
3. Choose the name of the file and the folder to which you want to export the file.

10.3 Use case: Managing access to files on a USB key

Access to files on a USB key can be blocked at several levels. SES Evolution verifies in this order:
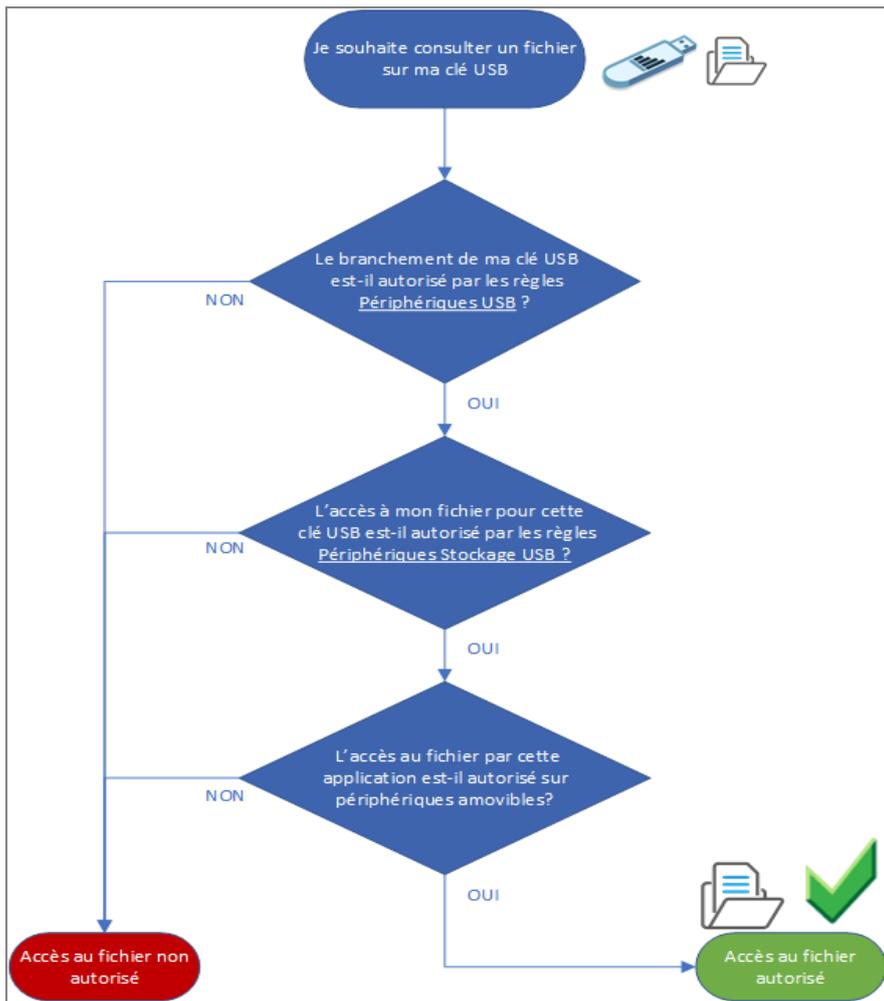
USB device rules:

- Verification of the USB key: USB class and sub-class, vendor ID, product ID and serial number.

USB storage device rules

- Verification of the path and file name,
- Verification of the USB key's vendor, product and serial number,
- Verification of the USB key's trustworthiness.

File/application ID ACL resource rules

- Verification of the file's or application's accessibility when it is located on a removable device.

Diagram text:
- Je souhaite consulter un fichier sur ma clé USB
- Le branchement de ma clé USB est-il autorisé par les règles Périphériques USB ? — NON / OUI
- L'accès à mon fichier pour cette clé USB est-il autorisé par les règles Périphériques Stockage USB ? — NON / OUI
- L'accès au fichier par cette application est-il autorisé sur périphériques amovibles? — NON / OUI
- Accès au fichier non autorisé
- Accès au fichier autorisé

10.4 Use case: Blocking access to USB keys that have not been decontaminated

Many malware programs can be spread through USB keys. To safely monitor USB keys plugged into your pool, you can make it mandatory to decontaminate all keys with contents that were modified outside the organization. To do so, set up air-gapped workstations equipped with antivirus solutions that analyze the plugged in devices. Next, configure SES Evolution so that it automates this analysis and guarantees that only USB keys with the appropriate level of trust are allowed on SES Evolution agents.

USB keys that are modified on a SES Evolution-protected workstation keep their trust level and do not need to be decontaminated.

10.4.1 Creating an agent group for air-gapped workstations

1. Create a *Decontamination* agent group of all the workstations used as USB key decontamination airlocks. For further information, refer to the section Creating agent groups.

2. Enable the following options in the **Trusted devices** section:

- **Allow device identification - Automatic**,
- **Automatically scan devices.**
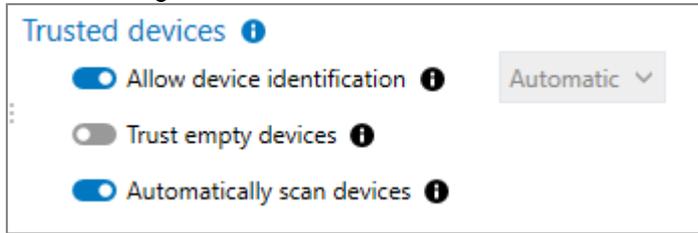
For further information, refer to the section Detecting and configuring the trust level on devices

10.4.2 Blocking USB keys based on their trust level

1. Create a **USB storage** security rule.
2. In the left section of the rule, add an *All* file ID that corresponds to all files.
3. Block all access in the default behavior.
4. Add a specific behavior that grants full access to keys with Trust level 2.

With this rule, full access can be granted to Trust level 2 devices, and those with a lower level will be blocked.

5. Apply this rule to all agent groups if you want to monitor the trust level of their USB keys.

For further information, refer to the section Controlling storage on USB devices.

11. Monitoring SES Evolution agent activity

SES Evolution offers an accurate view of SES Evolution agent and console activity through various types of logs classified by severity.

Among other data, logs contain the time of an event, the agent on which it occurred, the identity of the process that performed the operation, and if operations are blocked, information about the block.

11.1 Requirements

No short file names in MS-DOS 8.3 format must be visible in SES Evolution logs. Windows short file name creation must be disabled on all SES Evolution agents.

- To do so, set the value of the *NtfsDisable8dot3NameCreation* registry key to 1 in *HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\FileSystem*.

11.2 Various log types

SES Evolution agents generate several types of logs:

- **Event logs** are simple logs without an attached context. They provide information, for example, on blocked user actions that are prohibited by security policies, which then makes it possible to audit certain operations, etc. Events fall under several types:

Protection events	Generated when operations are blocked or audited by a security rule. For example, the process *illegimate_process.exe* attempted to run the process *abused_process.exe*.
Self-protection events	Generated when suspicious events are detected on the Windows system that are not associated with a security rule. For example, the user attempted to delete a protected file.
Operational events	Generated when events relating to the global operation of SES Evolution are detected. For example, the agent applied a new policy.
External events	Generated when events relating to *External event forwarding* and *OSSEC* audit rules are detected.
Windows Defender events	Generated when Windows events relating to the *Virus and threat protection* feature are reported. These logs are displayed only when the security policy contains *Stormshield - Windows Defender event forwarding* rule sets.

- **Alert logs** indicate that an attack occurred. Such logs come with a context that makes it possible to analyze the events that led up to the malicious action.
- **Context logs** are captured continuously on agents and represent an overall audit of actions performed on a workstation. They are not kept and are sent only when an alert is detected. These logs provide information on activity on the workstation just before and after the attack.

Agent logs can be read on the administration console and the agent's interface. They can also be read on the Syslog server if you have configured one.

Depending on whether you want to make changes or only view the **Agent logs** panel, you must have the **Agent logs-Modify** or **Agent logs-Display** privilege.

You can configure the log levels to send to the console, the agent and the Syslog server. For further information, refer to the sections Sending logs generated by agents and Configuring log management.

The agent has a protection mechanism against log flooding.

When it detects a certain number of strictly identical or similar logs over a short period of time, it stops generating the following similar logs and counts them. In addition, it does not generate a context even if the security rule associated with the log is configured accordingly. However, the protections remain active and the other logs are still generated.

It then issues a specific log indicating the detection of the log flooding. When log generation falls below a certain threshold, it issues another log to signal the end of the generation of similar logs. Depending on the log display setting, these two logs can be displayed on the agent interface and in the administration console.

In the administration console, from the logs indicating the start and end of the log flooding, you can access the log that triggered the protection. If necessary, create an exception on this log or adapt your security policies to prevent the phenomenon from recurring. To create an exception, see Adding exceptions to logs.

11.3 Viewing and managing agent logs in the administration console

All the logs you have configured to be displayed on the console are visible in the **Environment > Agent logs** menu. In this menu, you will be able to analyze, filter and manage logs, add exceptions so that certain logs will no longer be generated, and run Yara or IoC scans from logs. You can also Analyzing contexts to understand attacks and Managing remediation tasks from logs.

The date and time of the agent logs displayed on the console are based on the time zone set on the machine hosting the console.
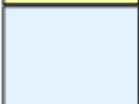
If an agent is offline and its logs were not sent to the agent handler, you can export its logs so that you can import and view them later in the **Agent logs** panel.

The **Agent logs - Modify** privilege is required to manage logs and create exceptions.

11.3.1 Reading logs

1. Select the **Environment > Agent Logs** menu.
   The list of logs for all components is displayed according to the active filters.
   When the logs panel is opened for the first time, the *New* and *In progress* logs, of *Emergency* and *Alert* level, issued in the last 24 hours are visible.
   Identical logs generated by several agents make up an *Event*. They are grouped by default
   on a single row and identified by the ⬚ icon.

2. If you wish to display the ungrouped logs view, disable the **Group events** option in the top right-hand corner.

3. To view details of logs in a group, click on the grouped logs icon, then expand the group by clicking on the + icon to the left of the group.

4. In the main log view or in the view of a log group, click on the **Date** button to select the period that you want to view, and click on **Apply**. With the double arrow in the drop-down menu, select the period from a calendar. The cross to the right of the **Date** field resets the period to the last 24 hours.
   The list of logs generated during the selected period appears.
   The color to the left of a log line indicates the severity level: there are 8 severity levels, corresponding to the Syslog protocol levels. Every level is assigned a color:
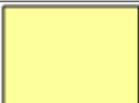
Urgence		Avertissement	
Alerte		Remarque	
Critique		Information	
Erreur		Diagnostic	

5. In the **Agent** column in the ungrouped view, click on the three dots to choose which information to display about the agent: Host name, user name and/or IP address.

6. To the right of the logs, click on ⬚ if you want to see or modify the rule that generated the log. This rule will stand out from other rules in the rule panel because it is grayed out and shows a blue bar on its left.

7. In the main log view or in a log group view, hover the cursor over the right-hand end of a log line, and click on ⬚ to open it and display additional information:

   - **Details** tab: full description of the processes, actions, etc. that caused the generation of the log. Links allow you to directly check:
     - The malicious nature of each process involved using the Google search engine or the VirusTotal website.
     - The reputation of a remote IP address on the VirusTotal and Stormshield IP Reputation sites.
     For this feature, the workstation that runs the administration console must have Internet access.

   - **Raw log** tab: code of the log in JSON format.

   If you suspect an issue and need to display more logs, change the log settings in the agent group logs or in the security rule.

11.3.2 Filtering logs

1. In the **Filters** table of the **Agent logs** panel, enable the filters to customize your list of logs. Every column corresponds to a type of filter and contains several values. Click these values to enable the corresponding filter, then click **Apply**.
   For example in this image, only the *New* logs of *Critical* severity are displayed.

- The number indicated in brackets refers to the number of separate logs, not the total number of logs. Identical logs make up a single event and are counted only once. The image above shows 9 critical logs (consolidated in two events), but only two logs are accounted for.

- The **Status** column allows you to filter logs by the status that you assigned. Refer to the section Managing logs.

- You can look for groups and agents in the **Agent group** and **Agent** columns by entering full or partial names in the search field.

- The **Application** and **Target application** columns make it possible to filter logs by applications that performed an action and those on which the action was applied.

2. Click **Advanced filters** to add other more precise filters and hence refine your list of logs. In the advanced filter window:

   a. Click on **Add filter**.

   b. Select the desired filter type. A line is displayed in the advanced filters window.

   c. Enter the value of the filter by selecting it from a list or by entering it manually.

   d. Specify whether the filter must include or exclude the value. This filter is inclusive by default – it displays all logs that match the chosen value. Click on [icon] to make this an exclusive filter.

   e. Add other filters if necessary. More advanced filters means fewer results in the list of logs.

   f. Click **Validate**.

   At any time, you can return to the initial filtering by clicking **Default filters**: only logs whose severity is *Emergency* or *Alert* and status is *New* or *In progress* will be displayed.

You can share a filter with a colleague who also has administrator privileges. For more information, see SES EvolutionSharing log information.

11.3.3 Managing logs

When you manage log analysis, you can assign a status to each log and indicate the name of the user who analyzed it. Log status is an important piece of information visible in the administration console dashboard in the **Events by status** area.

1. In the **Environment > Agent logs** panel, select one or several logs, then click on **Edit selected logs**. The **Edit logs** window appears.

2. In the **Status** list, select the status to allocate to the logs:

   - **New**: default status of a log. The log has never been analyzed.

   - **In progress**: the log is being analyzed.

   - **False positive**: the log has been identified as a false positive – a security rule triggered this log but it does not represent a malicious action. This status is automatically assigned to logs for which you have added an exception. For further information, refer to the section Viewing and managing agent logs in the administration console.

   - **Fixed**: the issue described in the log has been fixed.

   - **Closed**: the analysis of the log is complete. No further action is required.

3. In the **Assigned to** list, select a user name to assign the log to. This list shows all users declared in SES Evolution. For more information, see Managing users on the SES Evolution administration console.

4. In the **Comments** field, enter additional information if necessary about the log or your action. If anything is entered in this field, a tool tip will appear in the **Status** column in the list if logs.

5. Click **Validate**.

11.3.4 Performing a remediation from a log

If your appliance pool was targeted in an attack reflected in agent logs, you can perform a remediation to limit the impact of the attack and repair any damage caused.

1. In the **Environment > Agent logs** panel, select a log and click **Tasks > Create remediation task**.

2. Follow the procedure Managing remediation tasks.

11.3.5 Running a Yara or IoC scan from a log

If a log indicates that a process or file is potentially dangerous, you can configure a Yara or IoC scan to look up the process or file on the agent.

1. In the **Environment > Agent logs** panel, select a log and click **Tasks**.

2. Select the type of analysis.
   The task panel appears and the agent to which the log applies will be automatically selected.

3. Follow the procedure in Running a Yara scan on demand or Running IoC scans on demand to finish the configuration and run the scan.

11.3.6 Deleting events

In the **Agent Logs** panel, you can delete events that you no longer want to view. All logs contained in these events are then deleted from the log database. This action cannot be undone.

1. In the **Environment > Agent logs** panel, click **Actions > Delete events**, then select the desired action. The number of events affected by the action is displayed in brackets.
   - **Selected events** - to delete simple events or logs that you have selected manually in the list,
   - **Displayed events** - to delete all events and simple logs visible in the Agent Logs panel,
   - **All filtered events** - to delete all simple events and logs matching the filter criteria, including those that are not visible in the **Agent Logs** panel.

   The **Delete events** window appears.
2. Click **Estimate volume** to see the approximate volume of logs that will be deleted from the database.
3. Click on **Start**. A notification indicates that the operation is started.
4. In the notification, click **View manual deletion task** to go directly to the corresponding panel. For further information, see Deleting agent logs manually.

You can also automate log deletion on a daily or monthly basis. For further information, see Managing the deletion of logs.

11.3.7 Reading logs of offline agents

When an agent does not have access to the agent handler, its logs cannot be transmitted to the administration console and are therefore not visible in the **Agent logs** pane. You can export these logs to import them later into the console and read them in the same way as other logs.

Exported logs remain on the agent.

1. Log in to the agent workstation as an administrator.
2. Double-click on the 🛡️ icon in the status bar.
   The SES Evolution agent interface appears.
3. On the **Help and Support** 🧰 tab, click on **Assistance**.
4. Click the **Export unsynchronized logs** button and select the destination folder.
   A *cab* file is generated.
5. Copy it to a USB key or send it by email.
6. Copy this *cab* file to the import folder of the agent handler, e.g., *C:\ProgramData\Stormshield\SES Evolution\Server\AgentLogs\Import*. After approximately ten seconds, the file disappears from the *Import* folder and the logs it contains appear in the **Agent logs** panel.

> 💡 **TIP**
> You can also export the logs via a script, by running the EsGui program ([...]\Stormshield\SES Evolution\Agent\Bin\Gui) with the `/ExportLogs` command. Alternatively, you can specify in the command line the destination folder to which the file will be exported. For example: `EsGui /ExportLogs "C:\Users\Administrator\Desktop\Logs.cab"`

11.4 Adding exceptions for logs

If, after analyzing a log, you consider that the action that triggered it was not malicious and should not have been blocked, you can add an exception to the log. Doing so will prevent this action from being blocked and/or logged again in the future. Likewise, if you think that a file was wrongly quarantined, add an exception on the quarantine log.

Adding an exception creates a new exception rule in the exception rule set of your choice.

1. In the **Environment > Agent logs** panel, select one or more logs you no longer wish to generate.

2. Right click, then click **Add exceptions**. The **Create Exception Rules** wizard appears, allowing you to view the rule(s) that will be created by SES Evolution.

3. Expand the various boxes to display all the information on the rules. The policy(ies) affected by these rules are visible in the top right-hand corner.

4. Change the rules as you see fit:

    a. In the **Rule set** drop-down list, select the rule set in which the exception rules will be created: an existing *Exceptions*-type rule set, or a rule set created for the occasion with the label *NEW*.
    Click the *Pencil* button to edit the rule set name and type, or click the *Eye* button to view the rule set.

    b. Uncheck the rules or rule groups that you do not want to create.

    c. If necessary, change the rule intent, comment, application ID, or behavior for each action. For more information, see **Defining access control rules**

    d. Click the ⟲ button of a rule to reset its original values.

    > **ℹ️ NOTE**
    > If an identical exception rule already exists for the selected logs, it cannot be modified in the wizard, but you can modify it by clicking on the 🔲 button to navigate to the rule.

5. Click **Validate**.
   This will automatically:

   - Add one or more rules to the selected exception rule set. These rules ensure that a blockage does not occur in identical circumstances. The application or file IDs required for the rules are also created if necessary.

   - Assign a **False positive** status to the log in question, and identify the user who added the exception.

   - Add the comment *"Automatic exception created from this log"* to the log.

   - In an exception on a quarantined file, the file will automatically be restored at its original location the next time the **environment is deployed**.

6. If required, you can view or modify the exception rule created from the log:
   a. Display your log by enabling the **False positive** filter.
   b. Hover over the three dots to the right of the log line and select the **View exception rule** menu.

The exception rule that matches this log is displayed. It is distinguished from other rules by a blue bar on the left.
If the rule cannot be found, it has been deleted in the meantime.

11.4.1 Special cases

- If you request the creation of exceptions on several similar but not identical logs, the wizard displays an exception rule for each log:

In this case, you can change the ID of one of the rules to extend it to other similar rules using the wildcards * and ?.
In the example above, you can change the file ID to *c:\tmp\Log\Log** to include all similar logs in the same exception rule.
Other similar rules that match the new ID are automatically deselected in the wizard, and become false positives after validation.

- If the selected log is linked to a policy that no longer exists, you are notified and can choose whether or not to set the log to *False positive* status by activating the **Set event as false positive** button.

11.5 Sharing log information

To facilitate SES Evolution co-administration, you can share log elements with one or more colleagues. Sharing takes the form of a link, which is then copied by the colleague to their administration console. The shared elements are as follows:

- Log or log group,
- Attack chart,
- Agent log filtering.

11.5.1 Copying and sending the link

1. Select the **Environment > Agent Logs** menu.
2. Copy the desired link:

To copy the link of a ...	perform this action...
Simple log	1. Hover over the right-hand end of the log, group or context line. 2. Click ⋮, then select **Copy link**.
Log group	
Context log	
Detail view of a log	1. Open the detail view of a log. 2. Select the **Share > Copy link** menu in the top right corner.
Detailed view of a log group	1. Open the detailed view of a log group. 2. Click the `<` button in the top right corner.
Attack chart	1. Open an attack chart. 2. Select the **Share > Copy link** menu.
Log filtering	1. Apply one or more filters to the logs. 2. Click the `<` button in the top right corner.

3. Send the link to your colleagues, for example via email or instant messaging.

11.5.2 Displaying shared log items

You must possess **Agent Logs - View** privileges to display shared log items.

1. Retrieve the shared link, for example via email or instant messaging.
2. In the console's top banner, click on .
3. Paste the link into the field and click on the arrow to validate.
   The shared item is displayed.

11.6 Viewing logs in the agents' interface

1. On the workstation, click on in the status bar.
   The agent interface appears.

2. In the **Help & Support** tab 🧰, click **Events**.
   The list of logs from this workstation appears. An administrator user can view all log severity levels, while only *Alert* and *Emergency* level logs that have resulted in a block are displayed for a non-administrator user.
   The color on the left in a line of logs indicates its severity.
   The various color labels indicate:

   - The severity of the log, (e.g., Alert, Notice, etc.),
   - The type of log, (e.g., Internal, Self-protection, etc.),
   - The implemented protection, (e.g., Registry, etc.),
   - The action that SES Evolution applied (e.g., Block, etc.).

3. By default, you see only the logs accessible to the user who opened the session. Click on **Show all logs (administrator only)** to also see logs accessible to administrators. For example, if several users connect to the same workstation, you can view logs for all sessions with this option.

4. Filter the list of logs to show only those that are relevant to you:

   - Click on one of the labels of a log to show only the list of logs that have this label. For example, click on *Registry* to display all logs relating to this registry base.
     Active filters appear at the top of the window. Delete all filters to display all logs again.
   - In the **Search** field, enter one or several character strings and press Enter to show only logs that contain these strings.

If you suspect an issue and need to display even more logs, change the log settings in the agent group logs or in the security rule.

SES Evolution keeps 500 MB of log history. When this capacity is reached, the oldest logs will be deleted, beginning with logs of the lowest priority.

11.7 Sending agent logs alerts by email

You can configure SES Evolution to send e-mail alerts to recipients of your choice. Alerts are triggered by some logs generated on SES Evolution agents.

You must first configure an SMTP server. For more information, refer to Configuring an SMTP server.

You must hold the **Email Notifications-Modify** permission to configure the sending of alerts.

To send email alerts:

1. In the **Backoffice > System** menu in the administration console, go to the **Email Notifications** tab.
2. Click on **Edit** in the upper banner.
3. In **Agent log alerts**, click **Add rule**.
   The rule creation wizard opens.

4. Enter the rule settings:
   - **Rule name**.
   - **Prefix of the subject of the email** received by the recipient. By default, the subject of the email begins with *SES EVOLUTION*. The prefix allows you to apply a specific process to SES Evolution alert emails in your mailbox.
   - **Frequency** with which you wish to send emails containing alerts, from one minute to 24 hours. Emails will be sent when the agent reconnects to the agent handler.
   - **Log attributes** for which you wish to trigger alerts.
   - **Log severity levels** for which you wish to trigger alerts.
5. Click on **Next**.
6. In the field at the bottom of the screen, enter the email address of the user who will receive the alerts, select the language, then click on **Add**.
7. Add more email addresses if you wish to send alerts to several recipients.
8. Click on **Create**.
   The rule will be added to the table in **Agent logs alerts**.
9. Add other rules if necessary.

During the specified duration, if agents generate logs that match the rules, an email alert will be sent.

You can disable or enable rules again by clicking on the checkbox in the **Enabled** column. The action buttons to the right of a rule can be used to duplicate or delete it.

You can temporarily stop emails being sent by disabling the **Enable notifications** option.

SES Evolution can also be used to email alerts on system logs or all dashboard content. For more information, see Sending system log alerts by email and Sending dashboard indicators by email.

11.8 Analyzing contexts to understand attacks

Contexts in SES Evolution make it possible to thoroughly analyze the environment in which attacks occur on agents, and determine what these attacks consist of, where they come from and how they strike. To get this feature, your security policy must contain the built-in rule set *Stormshield - Audits of attack contexts*. For more information, refer to the section Understanding built-in rule sets.

> ✏️ **EXAMPLE**
> If the Execution flow hijacking protection mode blocks a malware program, analyzing the context will reveal which file caused the malware to launch, and where the file came from.

11.8.1 Understanding what makes up a context

Contexts consist of two components:

- The **simple details** show only alerts and logs regarding the creation and killing of all processes that were run on the agent within the attack perimeter. The simple details are shown by default in the detailed context report.

- The **full details** show all the logs that the agent produced in the attack perimeter, including those that do not usually appear in the administration console. For example, even logs that remained local on the agent or that were sent to a syslog server can be seen in the full details. They are generated by the *Stormshield - Audits of attack contexts* rule set of the default policy.
Depending on the agent group configuration, the display of the full details may need to be manually enabled.

11.8.2 Configuring contexts

- All *Emergency* and *Alert* logs agent logs are automatically contexts. In addition, some protections systematically generate contexts during an attack. This is especially the case for process hollowing, execution flow hijacking and heap spraying, among others. Some protection rules are also configured by default to generate contexts when actions are blocked, or even during suspected attacks that are not severe enough to be blocked. For more information, see Managing vulnerability exploitation and Defining access control rules.

- In the context details, the size, perimeter, type and frequency of reporting to the agent handler can be configured for each individual agent group. For further information, refer to the section Configuring context details generated by agents.

- You can also define the level of context detail to be sent to the Syslog server. For further information, refer to the section Creating groups of agent handlers.

11.8.3 Analyzing contexts to understand attacks

1. Choose the **Environment > Agent Logs** menu.
The full list of logs from all agents appears.

2. Click on the small arrow to the left of a context to open it. It is accompanied by the eye icon, ⊙ or ⧉. Even if grouped logs are shown, the context will contain only the log line with the highest severity. For further information on how to read logs, refer to the section Viewing and managing agent logs in the administration console.

   > **ℹ NOTE:**
   > You can also open external contexts exported earlier. Refer to Exporting contexts and viewing external contexts.

3. Click on the eye icon ⊙ to the right of the context to display the detailed view of all the logs that make up the context. This view consists of several sections:

   - **Context chart**: represents the attack launched on the agent in the form of a graph. It shows all the processes involved in the context and how processes are linked to one another.

   - **Context detail logs**: lists all the logs of events surrounding the attack. A filter is enabled by default and only alerts are shown. Change the filters according to your preferences.

   - **Details** or **Raw logs** view: additional information about the item selected in the graph. Raw logs are generated in JSON format.

   - **Remediation** view: makes it possible to select and launch the desired remediation operations.

4. When the view is opened, the context chart highlights with a small blue shield the item that was attacked. Click on the processes that come before it (i.e., parent processes) and read the related information in the right pane. The **Hash** is particularly useful in checking whether this process was already identified as malicious in the database of known malware.
A red struck-through seal on a process means that it was not signed by a digital signature certificate when it was compiled.

> ✏️ **EXAMPLE**
> In our example, several indicators show that the first process is suspicious:
> - There is a red seal on its icon, meaning that the process is not signed,
> - Its **Name** was randomly generated,
> - It was executed by Winword, which does not usually execute such processes,
> - Its **Path** *C:\Users\abott\AppData\Local\Temp* shows that it was run in a temporary folder.

5. Depending on the agent group configuration, the context details may not appear automatically. If you need additional information, click on **Request more details** so that the agent will report all information to the agent handler. For further information on configuration, refer to the section Configuring context details generated by agents.

6. To search for context logs, enter your character string in the **Search** field. The search syntax is as follows:

Aide

Généralités

La recherche s'effectue dans le type d'événement, dans le message et dans le log brut.

- La recherche n'est pas sensible à la casse
- L'espace (' ') est considéré comme un opérateur 'ET' implicite
- Le tiret ('-') permet d'exclure un mot-clé
- Les guillemets ('" "') autour des mots-clés permettent la prise en compte des espaces

Champs JSON

Vous pouvez rechercher des champs JSON dans le log brut en fonction de leur valeur.

- Expression : propriété_json **[opérateur]** valeur_json
- Le point ('.') permet d'indiquer une propriété imbriquée
- Les opérateurs disponibles sont :

Opérateur	Symbole	S'applique
Égal	=	À tous les types de caractères
Contient	%	À tous les types de caractères
Supérieur à	>	Aux caractères numériques
Inférieur à	<	Aux caractères numériques
Supérieur ou égal à	>=	Aux caractères numériques
Inférieur ou égal à	=<	Aux caractères numériques

Example

```
explorer.exe -"exécution de processus" type>=11 severity<4
createdprocess.processguid=539FE70B-6B8B-449B-98C9-0520366C5362
```

Searches will cover context detail logs.
Only logs that match the search will remain displayed in the list. Searches have no impact on the context chart.

> ✏️ **EXAMPLE**
> In our example, the command line of the *WINWORD.exe* process indicates that a file
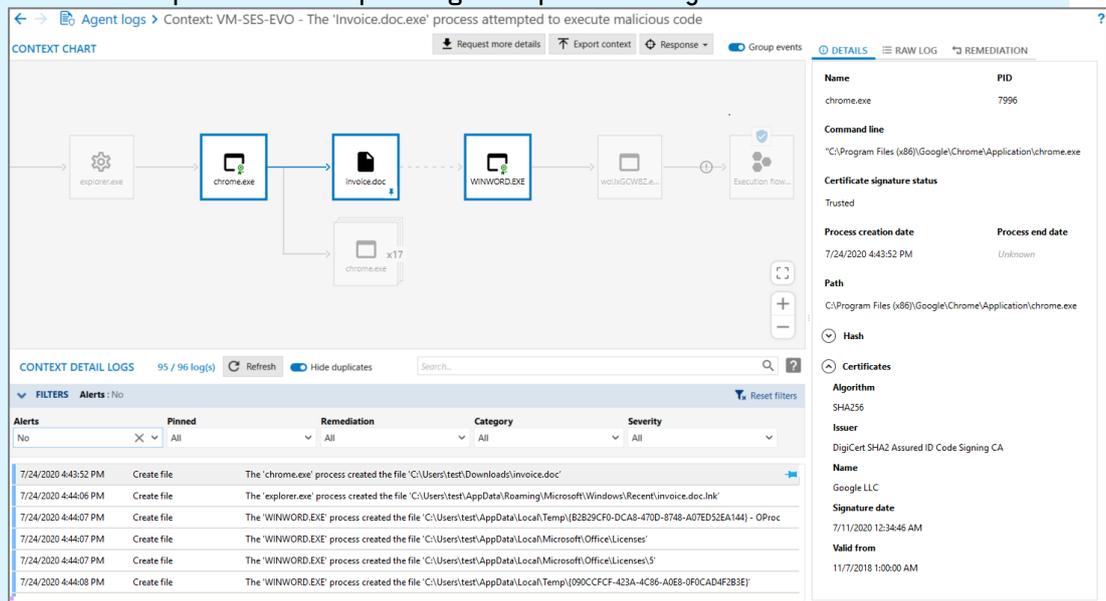
nvoice.doc was created. Searching for the string *invoice.doc "file creation"* will display all logs that include these terms and also reveal that *chrome.exe* created this file.

7. If you have identified a log that may help you understand the attack, pin it to the chart by clicking on 📌. This log will be added to the chart as a new event, and modifies the chart as a result.
To list only logs that match items in the chart, click on **Pinned only**.

> ✏️ **EXAMPLE**
> In our example, if you pin the log that mentions the creation of the *invoice.doc* file, you can understand how the attack was performed: the malware was launched on the workstation by an infected Word document (*invoice.doc*) that the user downloaded via Chrome and opened. It was a phishing attempt blocked by SES Evolution.

8. To examine a specific part of the chart more closely, move your mouse and zoom using the buttons at the bottom right of the chart. You can also use the left button on the mouse together with the scroll wheel.

9. Since identical processes are grouped by default, disable **Group events** at the top on the right to deploy items and analyze them individually.

10. Once the scan is complete:
    - In the **Remediation** tab at the top right side, create a remediation task by selecting the operations that you want to perform on the agents in question. Click on **See remediation task** to run the task. For more information, see the section Managing remediation tasks.
    - Click on the back arrow at the top on the left to return to the standard log panel. All your changes will be saved and appear the next time you open the contexts view.

11.8.4 Exporting contexts and viewing external contexts

Contexts can be exported, allowing you to:
- Share them with an external service for analysis,
- Archive them in a storage space so that they can be deleted later from the log database.

Once they are exported, these contexts can be viewed in an SES Evolution administration console.

Exporting contexts

1. Choose the **Environment > Agent Logs** menu.
2. The full list of logs from all agents appears.
3. Select a Context log. It is accompanied by the eye icon, [👁] or [🗔].
4. In the button bar at the top, click on **Contexts > Export contexts**.
5. The file is exported by default to the Desktop on the local workstation, and the name of the file is made up of the date and time followed by the name of the context. Change the name and location as needed.
6. If necessary, add a description in the **Comments** field, then click on **Export**.
   A notification will appear when the export is complete, allowing you to open the exported .cab file. The archive may contain up to four files: *contents.json*, *package.json*, *minicontext.txt* and *fullcontext.json*, the last being optional.

You can select several contexts to be exported simultaneously. In this case, their names cannot be changed, and you cannot enter comments.

Contexts can also be exported from the context chart.

Viewing exported contexts

1. Choose the **Environment > Agent Logs** menu.
2. In the button bar at the top, click on **Contexts > Open external context**.
3. Select the *.cab* file corresponding to the context you wish to view.

> 💡 **TIP**
> You can also drag and drop the *.cab* file into the blue area under the menu on the left side of the administration console.

The context chart then opens. You can use it in the same way as a context generated on your pool. For further information, refer to the section Analyzing contexts to understand attacks.
If you quit the chart, data will be kept for two minutes, and you can return to it by clicking on the **Environment > Agent logs** menu again. When the chart expires, you can open it to view it again.

12. Analyzing behavior on user workstations

SES Evolution makes it possible to run scans on user workstations, in particular searches for binary or textual schemas with the Yara tool, and searches for indicators of compromise (IoC). These scans, which can run as background tasks or are triggered when a particular event occurs, make it possible to detect malicious or suspicious behavior. This allows you to react quickly when a proven attack occurs.

Quarantined files are excluded from scans.

To ensure operation of these scan operations, the **Yara scan** and **IoC scan** features must be enabled in the agent group configuration. For more information, see Scheduling Yara scans and Scheduling IoC scans.

12.1 Running Yara scans

Yara is a tool that helps to identify and classify malware through rules. By using Yara rules, you can detect textual or binary patterns in files or processes being run on SES Evolution agents. In concrete terms, the integration of Yara scans in SES Evolution allows malware blocked by SES Evolution to be named, identified on other workstations and possibly quarantined or terminated.

> ⚠️ **CAUTION**
> Even though SES Evolution has been designed to limit their impact on workstations, Yara scans may still affect the performance of scanned agents. The scan time depends on the number of rules and their nature, but also the number of files to scan. We recommend targeting the directories to scan.
> The scan can also increase processor usage, and hence trigger the fans or use more battery power on a laptop PC.
> For more information, refer to the section Choosing the priority of Yara and IoC analyses.

To run Yara scans in SES Evolution, you must first import these rules into analysis units. The scan can then be run on agents in three different ways, as described in the scenario below.

> ✏️ **EXAMPLE**
> A malicious file *Invoice.doc* is emailed to all your company's employees. Some of them download and open it. When the file is opened, it runs a process on the workstation that performs malicious operations. By running Yara scans, the security administrator can:
> - Determine whether the process blocked by the *Protect-Office apps-Part 7* rule in the *Protection baseline* rule set is malicious. To do so, the administrator configures the rule so that when it is applied, it triggers a Yara scan. See Triggering a Yara scan when logs are generated in a rule.
> - Detect and shut down the malicious process on agent groups that are not protected by this policy. To do so, the administrator runs a scan on demand on the relevant agent groups. See Running Yara scans on demand.
> - Check daily for the presence of the *Invoice.doc* file on workstations in order to receive alerts. To do so, the administrator must configure a scheduled scan. See Scheduling Yara scans.

12.1.1 Getting Yara rules

To launch Yara scans on SES Evolution agents, you must have Yara rules. There are several ways to obtain some:

- Retrieve Yara rules shared publicly by organizations such as the CERT-FR.
- Download Yara analysis units available on the Stormshield server. These units contain Yara rules. For more information, refer to the section Downloading Stormshield updates.
- Create your own Yara rules. For more information, refer to the Yara documentation

12.1.2 Creating Yara analysis units

Yara analysis units consist of one or several Yara rules.

You must hold the **Resources-Modify** privilege to create analysis units.

Finding out Yara / SES Evolution agent compatibility

SES Evolution agent version	Compatible Yara versions
2.7	4.5.0
2.6	4.5.0
2.5	4.2.3
2.4	4.2.3
2.3	4.2.2 and 4.2.3

Creating analysis units

1. Select the **Security > Resources** menu.
2. In the left panel, click on **+ Add a resource**.
3. Select **Yara scan**, then the scan mode.
   - **File scan** to analyze files contained on agents,
   - **Process scan** to analyze memory on the processes that are executed on agents. Do note that this scan mode may detect the same pattern in several processes as memory from one process can be temporarily copied into other processes.
   The new unit will be added to the **YARA** category in the panel on the left. You will find the resources provided by Stormshield under the category **Stormshield YARA**.
4. In the **New analysis unit** field, enter the name of your analysis, then a description below it if necessary.

5. Click on **Import files** and select the **.yar*, and **.rule* files that you wish to use in this analysis unit. You can also import Yara Index files that reference other Yara files.

If imported Yara files contain inconsistencies or are likely to affect performance, Error, Warning or Performance messages appear in the **Compile resources** area with a description. If an error occurs, you must fix the issue or remove the file in question as you will not be able to save the analysis unit.
You can filter these messages by severity, Yara version or SES Evolution agent version if needed.

Yara analysis units cannot be deleted while they are being used in a Yara task, scheduled scan or as an action when logs are generated in an SES Evolution rule.
To obtain a local copy of the Yara files, if they were imported by another administrator, for example, click on ⌷ and select a destination folder. You can then look up these files, edit them and import them into the same analysis unit or into another unit.

You can also import a *.cab* file directly from the ▤ menu in the panel on the left. Cab files contain the file(s) to be used in a Yara or IoC analysis unit as well as other data such as the title and description of the unit. In the same menu, the **Export** sub-menu makes it possible to export an analysis unit with all this information in a *.cab* file.

You can also import a *.cab* file directly from the ▤ menu in the panel on the left. Cab files contain the file(s) to be used in a Yara or IoC analysis unit as well as other data such as the title and description of the unit. In the same menu, the **Export** sub-menu makes it possible to export an analysis unit with all this information in a *.cab* file.

12.1.3 Triggering a Yara scan when logs are generated in a rule

You can configure SES Evolution rules to automatically launch a Yara scan on an agent every time the rule is applied, i.e., every time a log is generated for such a rule. The types of rules in question are Threats, Application, ACL resources and Networks.

> ⚠ **WARNING**
> Yara scans triggered when logs are generated have a greater impact on the performance of agents than the impact of scheduled scans or scans on demand.

For more information on possible actions when logs generated, refer to the section Configuring actions triggered by rules.

12.1.4 Running Yara scans on demand

Unscheduled Yara scans can be run whenever needed. To do so, you must create a Yara scan task.

1. Select the **Responses > Manual tasks** menu and click on **Create a task**.
2. Select **Yara scan**.
   You can also open the tasks panel through **Agent logs** by selecting a log and clicking on **Tasks > Create a Yara scan task**.
3. Tick all the agents on which you wish to run the Yara scan. If required, use the filters to display only those agents meeting certain criteria, then click **Next**.
4. Give your task a **Name**.

5. Click on **Add scanunits** and select the scan units you wish to include in your Yara scan, then click on **Validate**.
Click on **Next**.

6. Click on Log settings to determine the severity and destination of the SES Evolution logs generated during the Yara scan.

7. In **File scan parameters**, select **Default scan** to run a recursive scan on the folder\\.\EsaRoots\SystemDrive and exclude the folders \\.\EsaRoots\SystemRoot, \\.\EsaRoots\ProgramFiles and \\.\EsaRoots\ProgramFilesX86. Otherwise, select **Custom scan**:

   - **Analyze the image file of running processes**: checks whether the *.exe* file in the processes contains the Yara pattern you are looking for. This option also allows you to shut down any malicious processes identified on agents during the Yara scan, and/or exclude from the scan any processes run by Windows administrator and/or system accounts.

   - **File extensions:** Restricts scans to the indicated extensions.

   - **Included files and folders:** runs the scan on indicated files and folders with or without recursion.

   - **Excluded files and folders:** excludes from the scan indicated files and folders with or without recursion. Click on the + icon to add another path.

8. In the **Process scan parameters**, select **Default scan** to run a memory scan of all the processes being executed on the workstation, otherwise, select **Custom scan**:

   - **Shut down the process detected**: Stops dangerous processes identified during the Yara scan.

   - **Exclude processes run by:** Excludes from the analysis the processes that were run with the indicated integrity levels (administrator and/or system).

   - **Directory of excluded processes**: Excludes from the analysis the processes for which the executable files are located in the indicated folders. Click on the + icon to add another path.
   You can also export scan settings in JSON format and import them again for other tasks.

9. Click on **Next**.

10. Click on **Run task**.
The task will appear in the main task panel.

11. Right-click on each task to perform the following operations:
   - Browse to the agent logs corresponding to this task,
   - Remove the task from the list,
   - Cancel the task currently being run on agents,
   - Run the task again by changing some settings.
   You can also **Delete completed tasks** from the tasks panel.

12. Click on the arrow to the left of the task to show details about the analysis units that the task contains.
Click on **Clear selection** to cancel a running analysis unit.

12.1.5 Scheduling Yara scans

Yara scans can be scheduled so that they can be run regularly by a group of agents.

For more information, refer to the section Scheduling Yara scans.

12.1.6 Looking up Yara scan usage

1. Select the **Security > Resources** menu.
2. In the left panel, click on the resource for which you want to view usage.
3. Deploy the **Usage of resource** area, which provides the following information:
   - **Agent groups - Scheduled scan**: Agent groups that were analyzed during a scheduled Yara scan for this resource.
   - **Rule sets - Action when logs are generated**: Rule sets that triggered the Yara scan. Click on the name of the rule set to open the corresponding panel.
   - **Tasks**: Tasks relating to Yara scans executed on demand. Click on the name of the task to open the corresponding panel.

12.2 Searching for indicators of compromise

IoC (Indicators of Compromise) scans make it possible to measure the extent of an incident or attack on a workstation by searching for indicators of compromise. Indicators may be, for example, malware signatures, specific IP addresses, malicious file hashes, suspicious URLs or text files. They can be searched for in DNS requests, Windows named objects or event logs for example.

Indicators may reveal the tools used and the perpetrators of the attack.

In order to search for indicators on users' workstations, you must first import lists of indicators into the analysis units in SES Evolution. The scan can then be triggered automatically when a security rule detects or blocks unusual behavior and generates a log. To protect workstations from potential attacks, you can also schedule the scan to run regularly and for a set period of time or run it on demand.

The logs generated from the IoC scans are then used to perform remediation actions to remove the detected malware. For more information, see the section **Managing remediation tasks**.

> ⚠ **CAUTION**
> Even though SES Evolution has been designed to limit their impact on workstations, IoC scans may still affect the performance of scanned agents. Scan time depends on the number of IoC indicators and their type, but also the number of files to scan. We recommend targeting the directories to scan.
> The scan can also increase processor usage, and hence trigger the fans or use more battery power on a laptop PC.
> For more information, refer to the section **Choosing the priority of Yara and IoC analyses**.

12.2.1 Creating IoC analysis units

In IoC analysis units, you can import the various types of indicators of compromise in the form of lists. You can then use these analysis units in security policy rules, manual on-demand tasks or scheduled tasks.

Refer to the following sections on how to use analysis units:

- **Triggering an IoC scan when logs are generated in a rule**
- **Running IoC scans on demand**,
- **Scheduling IoC scans**.

Requirements

- You must hold the **Resources-Modify** privilege to create analysis units.
- The indicators of compromise that you will use in analysis units must be compiled in CSV files. You can find examples of indicators on websites, such as the ANSSI's website. Indicators may also originate from your own device pool if you have detected a compromise with SES Evolution or via other means, or from Stormshield resources downloaded from the update server. For more information, refer to the section Downloading Stormshield updates. The CSV file may be in the following format: indicator in the first column, followed by a separator (comma, semi-colon or tab) and a description in the second column. The description is not mandatory, in which case, the separator will not be necessary. If the CSV file contains more than two columns, it can still be imported into the unit but only the first two columns will be taken into account and shown in the console. Column titles are not required.
  For hash indicators, the hash algorithm does not need to be indicated in the CSV file. The console will automatically detect it.

IoC hash example		
# IoC File hash : analysis of a file according to its content defined by a hash		

INDICATOR	HASH	DESCRIPTION
6b2904f8ee9f9b3e04ba263c7c6ade608c894c22...	SHA256	7zconfig.exe
6ef27bf59e66f5773bcfe20a35c97efc	MD5	scvhost.exe
15ff120b430021c36c232c99ef8d926aea2acd7b	SHA1	svjhost.exe
24:II9rFBzwjx5ZKvBBi8RuM4Pp6rG5Yg+q7wIXhI...	SSDEEP	svjhost.exe
24:MI9rFBzwjx5ZKvBBi8RuM4Pp6rG5Yg+q8wIXh...	SSDEEP	svjhost.exe
24:OI9rFBzwjx5ZKvBBi8Ru64Pp6rG5Yg+q8wIXh...	SSDEEP	svjhost.exe

Creating IoC analysis units

1. Select the **Security > Resources** menu.
2. In the left panel, click on **+ Add a resource**.
3. Select **IoC scan**, then the type of indicators to search for in the scan:
   - Text: search for suspicious character strings (file name, domain name, host name, IP address, subject name, email address),
   - File name: searches for suspicious files,
   - File hash: searches for suspicious file hashes, If you import SSDEEP file hashes, you can change the default similarity rate of 80% by clicking on **Global settings** in the top right corner of the **Resources** panel. A 100% match between two files means that they are identical.
   - DNS request: searches for requests submitted to suspicious domain names,
   - Named object: the search will cover suspicious object names among Windows named objects (ALPC port, Event, Job, Mutant, Section, Semaphore, Timer, Mailslot, NamedPipe, etc.).

   The new unit will be added to the **IOC** category in the panel on the left. You will find the resources provided by Stormshield under the category **Stormshield IOC**.
4. In the **New analysis unit** field, enter the name of your analysis, then a description below it if necessary.

5. Click on **Import .csv file** and select a CSV file that lists the indicators matching the chosen type. The total number of indicators appears to the left of the **Search** field. You can import only one CSV file per unit.

6. Click on **Save**.

IoC analysis units cannot be deleted while they are being used in an IoC task, scheduled scan or as an action when logs are generated in an SES Evolution rule.

To obtain a local copy of the CSV file, if it was imported by another administrator, for example, click on ⬆ and select a destination folder. You can then look up this file, edit it and import it into the same analysis unit or into another unit.

You can also import a *.cab* file directly from the ☰ menu in the panel on the left. Cab files contain the file(s) to be used in a Yara or IoC analysis unit as well as other data such as the title and description of the unit. In the same menu, the **Export** sub-menu makes it possible to export an analysis unit with all this information in a *.cab* file.

12.2.2 Triggering an IoC scan when logs are generated in a rule

You can configure SES Evolution rules to automatically launch an IoC scan on an agent every time the rule is applied, i.e., every time a log is generated for such a rule. The types of rules in question are Threats, Application, ACL resources and Networks.

> ⚠ **CAUTION**
> IoC scans triggered when logs are generated have a greater impact on the performance of agents than the impact of scheduled scans or scans on demand.

For more information on possible actions when logs generated, refer to the section Configuring actions triggered by rules.

12.2.3 Running IoC scans on demand

Unscheduled IoC scans can be run whenever needed. To do so, you must create an IoC scan task.

1. Select the **Responses > Manual tasks** menu and click on **Create a task**.

2. Select **IoC scan**.
   You can also open the tasks panel through **Agent logs** by selecting a log and clicking on **Tasks > Create an IoC scan task**.

3. Select all the agents on which you want to run the the IoC scan. Use filters where necessary to display only agents that meet certain criteria. Click on **Next**.

4. Give your task a name.

5. Click on **Add analysis units** and select the analysis units that you want to include in your IoC scan. Click on **Next**.

6. Click on Log settings to determine the severity and destination of the SES Evolution logs generated during the IoC scan.

7. For Text indicators, you can disable the IoC scan in files, processes or event logs by unselecting the **Text search** checkboxes.

8. In **File scan parameters**, select **Default scan** to run a recursive scan on the folder\\.\EsaRoots\SystemDrive and exclude the folders \\.\EsaRoots\SystemRoot, \\.\EsaRoots\ProgramFiles and \\.\EsaRoots\ProgramFilesX86. Otherwise, select **Custom scan**:

   - **Analyze the image file of running processes**: checks whether the .exe file in the processes contains the indicators you are looking for. This option also allows you to shut down any malicious processes identified on agents during the IoC scan, and/or exclude from the scan any processes run by Windows administrator and/or system accounts.

   - **File extensions:** Restricts scans to the indicated extensions.

   - **Included files and folders:** runs the scan on indicated files and folders with or without recursion.

   - **Excluded files and folders:** excludes from the scan indicated files and folders with or without recursion. Click on the + icon to add another path.

9. In the **Process scan parameters**, select **Default scan** to run a memory scan of all the processes being executed on the workstation, otherwise, select **Custom scan**:

   - **Shut down the process detected**: Stops dangerous processes identified during the IoC scan.

   - **Exclude processes run by:** Excludes from the analysis the processes that were run with the indicated integrity levels (administrator and/or system).

   - **Directory of excluded processes**: Excludes from the analysis the processes for which the executable files are located in the indicated folders. Click on the + icon to add another path.

10. In the **Event logs** section, select the types of logs to scan and from which date.

11. In the **DNS request parameter** section, indicate the date from which you want to analyze DNS requests.

12. Click on **Run task**.
    The task will appear in the main task panel.

13. For each task, click on the icons below to perform several operations:

	In the agent logs panel, displays logs corresponding to this task.
	Removes tasks from the list.
	Cancels the task currently being run on agents.
	Run the task again by changing some settings.

You can also **Delete completed tasks** from the tasks panel.

14. Click on the arrow to the left of the task to show details about the analysis units that the task contains.
Click on **Clear selection** to cancel a running analysis unit.

12.2.4 Scheduling IoC scans

IoC scans can be scheduled so that they can be run regularly on a group of agents.

For more information, see Scheduling IoC scans.

12.2.5 Looking up IoC scan usage

1. Select the **Security > Resources** menu.
2. In the left panel, click on the resource for which you want to view usage.
3. Deploy the **Usage of resource** area, which provides the following information:
   - **Agent groups - Scheduled scan**: agent groups that were analyzed during a scheduled IoC scan for this resource.
   - **Rule sets - Action when logs are generated**: rule sets that triggered the IoC scan. Click on the name of the rule set to open the corresponding panel.
   - **Tasks**: Tasks relating to IoC scans executed on demand. Click on the name of the task to open the corresponding panel.

12.3 Choosing the priority of Yara and IoC analyses

For on-demand or scheduled Yara and IoC analysis types, SES Evolution allows you to set the priority of the process in relation to other processes running on the user's workstation. You can choose between a low or normal priority.

1. In **Security > Resources**, click on **Global settings** to the right of the panel.
2. In **Scan priority level** select a value for each type:
   - **Low**: the analysis will run as a background task with little impact on the workstation's performance,
   - **Normal**: the analysis will be faster and may slow down the workstation's performance.

GLOBAL SETTINGS ✕

Choose the Yara and IOC default settings:

IOC file hash analysis - SSDeep setting ⓘ

Similarity rate (0 to 100 %) `80`

Scan priority level ⓘ

Yara scheduled analysis	Low ⌄
Yara task	Low ⌄
IOC scheduled analysis	Low ⌄
IOC task	Low ⌄

✓ Save   ✕ Cancel

13. Responding to security events

When a malicious operation occurs or is suspected to have occurred on your pool, SES Evolution can detect and/or block it while generating a security event. It is also possible to respond to the event by performing a remediation on affected workstations. Remediation is a set of operations that make it possible to limit the impact of attacks and fix any damage caused.

> ✎ **EXAMPLES**
>
> - An SES Evolution audit rule monitors certain trees in the registry base to detect the addition or modification of any keys or values. This is crucial as some malicious programs use this method to persist after the workstation is restarted. If such an action is detected, an audit log will be generated, and you can launch a remediation to automatically delete or modify suspicious registry keys on the affected workstations.
> - A ransomware program was able to encrypt several files before SES Evolution blocked it. Remediation makes it possible to automatically retrieve the unencrypted version of these files from a Windows shadow copy.
> - A user unintentionally launched a malicious program. SES Evolution blocked it and prevented it from being run, but it can also be quarantined in addition. It is thus out of the user's reach, allowing the administrator to analyze it before deleting or restoring it.
> - Agent logs report suspicious events on a user's workstation. You detect a danger. You can isolate that workstation from the network in order to perform your analysis, which prevents an ongoing attack from spreading to the entire pool.
> - After a ransomware attack blocked by SES Evolution, some programs can continue to run on the workstation and facilitate new attacks, for example a remote access Trojan (RAT). You can detect such programs with an IoC scan, then delete them automatically with remediation.

13.1 Managing remediation tasks

You can run remediation tasks on workstations from agent logs. Depending on the type of logs, SES Evolution offers various remediation actions, such as file quarantining, registry key deletion, process shutdown, etc., which make it possible to greatly limit the impact of attacks.

The remediation actions requested may override the security policy in force on the workstation concerned.

13.1.1 Granting remediation permissions

Two separate permissions make it possible to manage access to the remediation feature. The **Remediation (advanced) - Modify** permission makes it possible to perform all remediation operations, including the ability to run Powershell scripts.

Since the execution of Powershell scripts is a very sensitive operation, grant this permission only to a very small number of trusted users.

For others who may perform remediation operations without running scripts, grant them the **Remediation - Modify** permission.

See Managing users on the SES Evolution administration console.

13.1.2 Creating a remediation task

1. Select the **Environment > Agent logs** menu and identify the log corresponding to the malicious operation that you wish to remediate. For example, a ransomware attack generates the log "The `process_name` process attempted to run a ransomware attack." All agent logs can be used to create a remediation task, except those with the attributes *Internal*, *Auto-protection*, or *Remediation*. For further information, refer to the section **Viewing and managing agent logs in the administration console**.

2. Select the relevant log or log group and click on **Actions** > **Response** > **Create remediation task**. The task window appears. It lists the possible remediation operations for the selected log type and the resource in question. Operations are grouped by agent.

3. Enter a **Name** for your remediation task.

4. Use filters to view only a certain **Action type**, or the actions concerning a certain **Agent group**.

5. Select the actions that you wish to perform. Depending on the selected log type, the following actions are possible:

   • Shut down a process, including or excluding its child processes,

   • Remove files (quarantine or delete);

   • Remove a folder and its contents (quarantine or delete);

   • Delete a registry key,

   • Remove or modify registry values;

   • Retrieve files encrypted by ransomware, with a list of the first 10 files encrypted;

   • Run a PowerShell script.

   For *Remove folder* and *Remove file* actions, you can choose to **Bypass exceptions**, *i.e.* perform the action even if the item is on the exception list.
   Some actions may contain an orange dot. This means that they affect one of the following critical system folders and may therefore have an impact on the workstation:

   • C:\Windows\System32

   • C:\Windows\SysWOW64

   • C:\Windows\Microsoft.NET

   • C:\Windows\WinSxS

   • C:\Program Files

   • C:\Program Files (x86)

   • C:\ProgramData\Stormshield

6. If the suggested actions are insufficient, and you hold the **Remediation (advanced) - Modify** permission, you can run a custom PowerShell script during the remediation task.

> ✎ **EXAMPLE**
> If a malicious program has added registry keys to persist after the workstation has been restarted, you may want to delete them. However, it can be tedious to select all the keys to be deleted if there are many of them and they affect several agents. It is therefore worth creating a script that automatically deletes all keys without the need to select them individually.

To add a Powershell script:

a. Click on **Powershell script actions > Add to all agents**.
The **Add a Powershell script action** window will appear.

b. To the right of the **Script** field, click on **+** to add the script to run.

c. In the **Arguments** field, specify if necessary the arguments to add when the script is run.

d. If you want the script to be run on all affected agents during the remediation task, check the **Select action for all agents** checkbox. Otherwise, the line will be added to the list of actions but will not be selected.

e. Select an existing script and click on 👁 to view it or ⬆ to import a new version of the script.

7. Click on **Start remediation**.
The panel for **Tasks** appears and you can track the progress of the remediation task.

8. Click on the arrow to the left of the task to show its progress on each affected agent.
The following are the various possible statuses of a remediation task:

Status	Description
Not started	The task was launched but has not yet started.
Running	The task is running.
Done	The remediation action was successful.
Error	An error occurred during the remediation task. A separate message indicates the reason. For example, the resource is locked, the user did not have sufficient privileges to delete files, the agent was not connected, etc.
Partial	• In a remediation following a ransomware attack, not all the files could be recovered. • If processes were deleted, at least one process could not be deleted.
Canceled	The task was canceled by the user while it was being run.

9. If you have chosen to quarantine files, they will be shown in the **Responses > Quarantine** panel. For more information, see the section Managing file quarantine.

13.1.3 Managing remediation tasks

Use the **Status**, **Type**, **Created by**, or **Agent group** filters to display only those tasks you want to see.

Click on **Details** for more information on the remediation actions performed and their results.

You can also perform the following actions on remediation tasks or sub-tasks:

- Cancel tasks in progress,
- Browse to the agent logs corresponding to this task,
- Run a task again,
- Remove a task from the task panel,
- Export the result file in CSV format.

13.2 Managing ransomware attacks

SES Evolution protects your organization's workstations from ransomware attacks. It can detect operations that ransomware applications usually perform on a system, such as file modification or encryption, and quickly stop them. If ransomware encrypted some files before being blocked by SES Evolution, you can retrieve the lost data by performing a remediation.

> ⚠️ **WARNING**
> SES Evolution's creation of shadow copies cannot replace regular backups. You must have a dedicated parallel backup solution.

13.2.1 Requirements

To block ransomware attacks, SES Evolution must be configured as follows:

- **Enable ransomware protection**. If you are using *Default policy* or *Backoffice component protection* policies, this protection mode is enabled by default in the Anti-ransomware protection rule set.
- **Enable Windows shadow copies**.
- Optional: Prohibit the execution of malicious commands that specially aim to delete shadow copies. To do so, use **application filtering via command line arguments**. If you are using *Default policy* or *Backoffice component protection* policies, this protection mode is enabled by default in the Anti-ransomware protection rule set.

13.2.2 Detecting ransomware attacks

If you have enabled SES Evolution anti-ransomware protection and chosen the **Block and kill** or **Block, kill and quarantine** option, every ransomware attack generates an Alert level log and a context:

"The process `process_name` *attempted to run a ransomware attack. Look up the list of encrypted files in the* `file_path` *file.*

The log contains:

- The name of the process behind the attack,
- The path of the remediation file that identifies all the files that the ransomware encrypted before it was blocked. This file will be kept for 30 days in *PROGRAMDATA%\Stormshield\SES Evolution\Agent\Diagnostics\Ransomware Protection* on the same workstation where the SES Evolution agent has been installed.
- The list of the first ten encrypted files (in the detailed log).

13.2.3 Retrieving lost data

If you have set up the requirements, you can retrieve lost data with the help of a remediation task, which allows you to retrieve an older version of the lost files.

Stormshield recommends that you retrieve the data 5 days later as SES Evolution will continue to create shadow copies daily after an attack. Since only the last five copies will be kept, new shadow copies containing encrypted files will overwrite older shadow copies.

To retrieve lost data:

1. Follow the procedure described in Managing remediation tasks.
2. When creating the task, select actions such as **Retrieve files encrypted by ransomware**.
3. Click on **Start remediation**.
   The **Manual tasks** panel appears and SES Evolution will start retrieving the encrypted files from the Windows shadow copies.
4. Once the task is complete, click on **Details** to view the restored files.
5. In Windows Explorer, check that the restored files have been saved in the original folder under their original names. The encrypted files will also be saved with a *.bak* extension.

13.3 Managing file quarantine

When a malicious operation occurs in your pool, SES Evolution makes it possible to detect suspicious files and quarantine them while they are being analyzed. Quarantined files can no longer be run, or cause any damage to the workstation. After the analysis, if the files are found to be harmless, you can restore them to their original location.

You can establish a list of folders to exclude. The files that they contain will be protected from quarantine.

Quarantine and restore operations are logged in Agent logs.

You need to hold the **Remediation-Modify** permission to quarantine and restore files.

13.3.1 Protecting files from quarantine

1. Select the **Responses > Quarantine** menu and click on the **Parameters** tab.
   In the **Predefined exclusions** section, some system folders are excluded by default, as it would be inappropriate to quarantine their content. Folders are displayed in the form of EsaRoots variables.
2. Click on **Edit** in the upper banner.
3. Under **Custom exclusions**, enter the path to the folder containing the content you wish to protect. Generic characters are not allowed.
4. Enable the **Recursive** option if the content of sub-folders needs to be protected as well.
5. In the **Owner** field, select one of the Windows groups that owns the file if necessary, or enter a specific SID.
6. Click on ➕ to validate the line.
7. Add as many paths as needed and click on **Save**.

You can temporarily disable exclusions by unselecting the checkboxes on the left side of the lines.

13.3.2 Quarantining files

Files can be quarantined:

- Automatically when a protection rule that you have configured is triggered. For more information on rule configuration, refer to Defining access control rules.
- Manually during remediation. For more information, see the section Managing remediation tasks.

Files located on a network share will not be quarantined.

During quarantine, files will be moved to the agent's local folder: *C:\ProgramData\Stormshield\SES Evolution\Agent\Quarantine*. Access to this folder is not allowed, even to administrators, and the files that it contains are encrypted.

13.3.3 Monitoring quarantined files

1. Select the **Responses > Quarantine** menu and the **General** tab to display the list of quarantined files.
2. If necessary, use the **Quarantine status** and **Agent group** filters to narrow down the list. The statuses can either be *Quarantined* or *Pending restoration*.
3. Select a file in the list to display information on the file and the agent concerned in the panel on the right.

13.3.4 Restoring quarantined files

After the analysis, if you consider the file harmless and a false positive:

- Add an exception on the log and deploy the changes on all agents, so that the file will no longer be detected as malicious,
- Restore it to its original location. The file will be restored exactly as it was, with the same ACLs and same alternate data streams.

To restore the file:

1. Select the **Responses > Quarantine** menu and the **General** tab.
2. Right-click on the file to restore and select **Restore selection**.
   The **Restore quarantined files** window appears. All quarantined files with the same hash will be listed and selected for restoration.
3. Enable the option **Overwrite existing file** if the same file already exists in the original location and you wish to replace it.
4. If necessary, use the search field or the **Quarantine status** and **Agent group** filters to narrow down the list. The statuses can either be *Quarantined* or *Pending restoration*.
5. Unselect any files that you would like to keep quarantined. All files with the same hash will always be restored at the same time at their respective locations.
6. Click on **Restore**.
   A restoration task will be created and the status of the files will switch to *Pending restoration*. The files will then be moved from the quarantine repository to their original locations, and disappear from the **Quarantine** panel.

13.3.5 Deleting quarantined files

Files are kept in the quarantine repository for 40 days before they are automatically deleted. In addition, if the volume of the quarantine repository exceeds 1 GB, the oldest files will be automatically deleted to make space for new files.

You can also choose to manually remove files from the **Quarantine** panel. In this case, the files will no longer be displayed, but will remain on disk in the quarantine repository. They will then be automatically deleted after 40 days, or if the 1 GB limit is exceeded.

1. Select the **Responses > Quarantine** menu and the **General** tab.
2. Right-click on the file you want to delete, and select **Delete selection**.
3. Confirm deletion.
   The file will no longer appear in the list.

Quarantined files will automatically be deleted when the SES Evolution agent is uninstalled.

13.4 Isolating computers from the network

When an attack occurs or is suspected in your pool, the affected computers can be isolated from the network. By Isolating computers, incoming and outgoing connections can be shut down immediately, preventing any attack from spreading to the rest of the network, or exfiltrating data to the attacker's servers.

While computers are being isolated, communication between agents and agent handlers is maintained, so that you can perform analysis and remediation operations if necessary. When the intervention is complete, you can undo the isolation of the computers and restore connections.

From the SES Evolution administration console you can:

- Isolate computers,
- Monitor isolated computers,
- Run Yara or IoC scan tasks and remediation tasks on isolated workstations,
- Undo isolation of the computers.

Requests to isolate and undo isolation are logged in System and Agent logs.

13.4.1 Requirements

The isolation feature can be used if all the following conditions are met:

- Computers can be isolated on SES Evolution agents as of version 2.5.3.

> **ⓘ NOTE**
> If you request the isolation of a set of agents in different software versions, only eligible agents will be isolated, i.e., those in at least version 2.5.3.

- The SES Evolution agent must be able to communicate with its agent handler to apply requests to isolate computers and remove them from isolation. Agents that are disconnected from the agent handler cannot apply such requests.
- The **Networks** feature must be enabled in agent group settings:
- You must hold the **Remediation-Modify** permission to isolate computers, undo their isolation and edit the list of connections allowed during isolation.
  If you hold the **Agent groups - Show** permission, you can look up the isolation status of computers in the agent group panel.

13.4.2 Isolating computers

Computers can be isolated regardless of the security policy applied to agents.

The feature can be accessed from two panels in the console:

- In the **Agents** panel, from the **All agents** view or from an agent group view, select one or more agents from the list and click on the **Isolate computers** button at the top of the panel.

- In the **Agent logs** panel, you can isolate one or more computers directly from a log or log group if you detect a suspicious event. Select one or more logs and click on the **Actions > Response > Isolate computer** button at the top of the panel.

The following window will appear, allowing you to create the isolation task:

The agents selected during the isolation request appear in the list. You can filter the display.

1. Enter a comment.

2. Check the selection of agents to be isolated. If your agent selection includes agents installed on server workstations, confirm that you wish to isolate these workstations. Application services installed on these workstations will no longer be accessible (web or mail server, file server, etc.).

3. Click on **Isolate**.
   The isolation monitoring panel appears (**Responses > Isolation** menu). For more information, refer to the section Monitoring and analyzing isolated computers.

As soon as the agent receives an isolation request:

- TCP and UDP connections that have already been opened to outside networks will be shut down. However, you can allow some connections to remain open. For further information, refer to the section Allowing network connections during isolation.

- New incoming and outgoing connections are no longer possible, except for connections necessary for communication with agent handlers (HTTPS port 443 for an agent greater than or equal to 2.7.1 or TCP port 17000 for an agent less than 2.7.1) and authorized connections. For further information, refer to the section Allowing network connections during isolation.

You can shut down isolated computers, as their isolated status will be kept when they are restarted.

> **ⓘ NOTE**
> You can move isolated agents from one agent group to another. They will retain their isolation status.
> Conversely, isolated agents automatically lose their isolation status in the following cases:
> - If you request the downgrade of the agent to a version below 2.5.3,
> - if you disable the **Networks** feature in the agent group settings.

13.4.3 Monitoring isolated computers

You can monitor isolated agents through the **Isolation** dashboard and panel.

The **Isolated agents** diagram in the dashboard shows the number of agents isolated from the network and their isolation status: Isolated, To be isolated, Undoing isolation pending. For further information, refer to the section Checking agent status.

You can perform the following actions in the **Responses > Isolation** panel:

- View the list of agents affected by the isolation.
- Undo isolation of the agents. For further information, refer to the section Undoing isolation.
- Create Yara and IoC scan tasks directly on the affected agents (**Threat hunting** button). For further information, refer to Analyzing behavior on user workstations.
- Determine the network connections allowed during isolation (**Settings** tab). For further information, refer to the section Allowing network connections during isolation.

13.4.4 Allowing network connections during isolation

By default, isolation blocks all network connections on the computer using TCP and UDP protocols, except for DNS and DHCP requests on ports 443 (agent version 2.7.1 or higher) or 17000 (agent version lower than 2.7.1) in order to allow communication with SES Evolution agent handlers.

The **Settings** tab in the **Isolation** panel makes it possible to allow other connections during isolation, by defining exception rules.

To add rules:

1. Click on **Edit** in the upper banner.
2. Click on **Add an authorized network connection**.
3. Enter a description in the window that appears.
4. Specify the path or SID of the application to be authorized.
5. Configure the settings that follow and click on **OK**.
   The rule will be created. It applies to all agents regardless of their agent group.

6. Click on **Save** in the upper banner.
7. Deploy the environment from the **Security > Deployment** menu.

All changes made in this tab will be logged in System logs.

This list of exceptions takes priority over the network rules in the security policy. It does not have priority over the agent's self-protection network rules.

13.4.5 Undo isolation

Any administrator with the necessary permissions can stop the isolation of a computer, even if another administrator isolated it.

As soon as the agent receives the request to undo isolation, TCP and UDP connections will be restored.

To stop isolating a computer, go to the **Isolation** panel, **General** tab:

1. On an agent's line or on a selection of agents, right-click > Undo isolation for the selection,
2. Enter a comment,
3. Check the selection of agents,
4. Click on **Undo isolation**.

13.4.6 Explanations on how isolation and challenges work

Maintenance mode

- Agents in maintenance mode can receive requests to isolate or undo isolation.
- The agent's maintenance mode can be enabled on computers that are under isolation. The computers will retain their isolation status.

For further information on Maintenance mode, refer to the section Enabling Maintenance mode.

Stopping the agent

When a computer is isolated and you request to stop the agent via a challenge:

- Network connections will be allowed again,
- The agent will still be seen as isolated in the administration console, and network connections will be shut down again at the end of the challenge.

When an agent is being stopped via challenge and you make an agent isolation request:

- The isolation request will be applied, but will not take effect until the challenge has ended.

For more information on stopping agents, refer to the section Stopping an agent.

13.4.7 Explanations regarding the maintenance of isolated agents

Automatic deletion of agents

The configuration for the automatic deletion of offline agents also applies to isolated agents. Agents are deleted by default after 30 consecutive days of staying offline.

For further information, refer to the section Monitoring agents in real time.

Uninstalling agents

Isolated agents can be uninstalled. In this case, if an agent is connected to its agent handler, the administration console will no longer consider it isolated from the network.

If it is disconnected from its agent handler when it is uninstalled, the management console will still consider it isolated. It will be automatically removed by the mechanism that automatically deletes offline agents.

13.4.8 Limitations of isolation

Starting up in safe mode

Network connections will not be shut down when an isolated computer is restarted in networked safe mode.

Computer disconnected from the corporate network

When a computer is no longer on the internal corporate network, as in the case of a mobile user for example, it cannot receive requests to isolate or undo isolation. The request will be applied when the computer is reconnected to the network.

Temporary web access

Temporary web access has a lower priority than network isolation.

When you request to isolate a computer with temporary web access :

- Temporary access will be shut down and the computer will be isolated,
- Temporary access will resume when isolation is undone, with the remaining time before isolation if the computer is not restarted during its isolation,
- Temporary access will not resume when isolation is undone if you restart the computer during isolation.

For more information on temporary web access, refer to Allowing temporary web access.

14. Downloading Stormshield updates

Stormshield regularly provides:

- New built-in security policies,
- Updates of built-in rule sets or new rule sets,
- Yara resources,
- IoC resources.

These resources can be found on the Stormshield public server or on a local server of your choice if you work in an offline environment. For more information on how to configure customized servers, see Configuring the Stormshield update server.

You can download them at any time in your administration console separately from SES Evolution updates.

When new policies, rule sets or Yara and IoC resources are available, an indicator will appear on the ⬛ icon in the upper banner of the console. Click on the icon to go to the panel in which updates are downloaded.

If you want to download resources that you do not already have, use the **Install all** button in the ≡ menu or the **Install** button in the panel to the right of each category. In this case, the new versions of resources that you already have will not be installed.

If you want to download new versions of resources that you have, use the **Update all** button in the ≡ menu or the **Update** button in the panel to the right of each category.

You must have at least the **Show** privilege on **Policies** or **Resources** to view the updates available in this panel.

Version release notes describe the new features in each resource. You can refer to them in a PDF file by clicking on ⬛ on the row of the resource, if the PDF file is on the update server.

After you download an update for rule sets that have already been deployed, a new version of the policies in question will be automatically created if you have selected **Always use latest version** of rule sets and there is nothing more you need to do.

After downloading Yara or IoC resources that are already used and deployed in rules, a new version of the rule sets that use these resources will be automatically created.

You can find the new Yara or IoC resources in the **Security** > **Resources** menu, under **StormshieldYARA** and **StormshieldIOC**.

You must hold **Edit** permissions on **Policies** in order to perform these operations.

Depending on user roles, you can make this panel inaccessible or display it in read-only mode, by using the **Updates** permission in the user panel. If you select **None** as the privilege, the panel in which updates are downloaded will not be visible.

For further information, refer to the section Managing users on the SES Evolution administration console.

To disable notifications when a new update is available, unselect **Enable notifications** at the top right side of the panel.

For more information on the resources available in this panel, refer to the sections Understanding built-in and custom security policies, Understanding built-in rule sets and Analyzing behavior on user workstations.

15. Managing backoffice components

The SES Evolution server consists of several backoffice components:

- One or several **agent handlers**,
- One or several **backends**,
- Two SQL Server databases: one for administration data and the other for logs,
- One or several **administration consoles**,

Logs generated by SES Evolution can be consulted when an event occurs on any of these components.

You can also monitor the size of the log database, configure an SMTP server to send e-mail alerts and indicators, configure the Stormshield update server, etc.

15.1 Monitoring the activity of SES Evolution backoffice components

The activity of backoffice components installed by SES Evolution generates logs the can be looked up in the administration console.

The **System logs - Show** privilege is required to read and filter logs.

To read system logs:

1. Select the **Backoffice > System logs** menu.
   Logs from all components appear according to the active filters applied. The first time the log panel is opened, the logs displayed will be all the logs that were generated over the past 24 hours.

2. Click on the **Date** button to select the period that you want to view, and click on **Apply**. With the double arrow, select the period from a calendar. The cross to the right of the **Date** field resets the period to the last 24 hours.
   The list of logs generated during the selected period appears.

3. In the **Filters** table, enable filters to customize your list of logs. Every column corresponds to a type of filter and contains several values. Click on these values to enable the corresponding filter.
You can look for hosts and users in the **Machine** and **User** columns by entering full or partial names in the search field.
You can go back to your initial filtering at any time, by clicking on **Default filters**: all logs will be displayed again.
The color on the left in a row of logs indicates its severity:

- Blue: Information,
- Yellow: Notice,
- Orange: Error.

4. Click on the small arrow to the left of the log to open it and display additional information:

- **Details** tab: Full description of the event that caused the generation of the log.
- **Raw log** tab: code of the log in JSON format.

15.2 Monitoring databases

SES Evolution uses an administration database and a log database. Depending on the configuration of your security policies and the desired amount of logs, the volume of these databases may grow quickly and reach a critical threshold.

With the calculation of estimates and alerts, SES Evolution allows you to monitor the remaining capacity in databases and anticipate when they will be saturated. You can also schedule maintenance sessions or when logs will be deleted.

15.2.1 Looking up general information on databases

1. Select the **Backoffice > System** menu and click on the tab of the database required, Administration or Logs.

2. In the **SQL Server** section, look up the following information:

Instance	Name of the SQL Server instance.
Version	SQL Server version used.
Last Connection	Number of seconds since the backend's last connection to the database.
Status	• **Green**: All databases can be reached and less than 71% of the disk space is used. Saturation of the database is estimated in more than three months. • **Orange**: On at least one database, between 71% and 80% of the disk space is used, or saturation is estimated in one to three months. • **Red**: At least one database cannot be reached, or more than 80% of the disk space is used. Saturation of one of the databases is estimated in less than a month.
Database files. View under the table the detail of the log database files.	• **Usage**: Percentage of the available space currently used by the file. • **Space used**: Space currently allocated by SQL Server for the file. • **Capacity**: Space available on the disk for the file.

The log database ██████ the following three files:

- EsLogs: contains the reference tables centralizing redundant data present in the logs; for example, the agents having issued the logs, the source and destination applications, the paths of the files and register keys. In particular, this information speeds up the search process.
- EsLogs_log: this is the SQL Server transaction log, which contains temporary technical data.
- EsLogs_Events which contains:
  - the logs of the agents, including the contexts associated with the events when they exist,
  - the system logs.

15.2.2 Monitoring database size

SES Evolution allows you to monitor the size of databases in several ways: through a chart, and through the automatic degraded mode for the log database.

Estimating and monitoring the size of databases through a chart

1. Select the **Backoffice > System** menu and click on the desired database tab, Administration or Logs.
2. In the section **Database size supervision and estimate**, refer to the chart. It shows information on the occupation of the data over 9 months. For more information, see the section Understanding the database tracking graph.
3. A *Warning* log appears by default in the system logs three months prior to the database's estimated date of saturation, and an *Error* log appears a month before. Only in the log database, if the default values do not suit you, click on **Edit** in the upper banner and select the desired number of months by changing the following settings:
   - **Generate a warning n months before estimated date of saturation**
   - **Generate an error n months before estimated date of saturation**

   A colored banner will also appear when the estimated date of saturation approaches: orange between one and three months from the date, or red from one month onwards.
4. Click on **Save** at the top right of the window to save changes.

Understanding the database tracking graph

The **Database size supervision and estimate** chart is divided into two parts:

- The left side shows six months of database occupation history up to the current date. The database size is measured every day at *00:00* UTC. The points on the chart correspond to the measurement taken on the first day of every week.
- The right side shows the projected increase in database occupation over three months from the current date. For greater reliability, it starts appearing only after the database has been used for 14 days.
  The first orange dot indicates the date on which the database will be saturated in a pessimistic forecast model. The second orange dot represents the same date in an optimistic scenario.
- **Disk space usage**: disk space used in relation to available space, and percentage of occupation.
- **Estimated saturation**: period in which the database is estimated to reach saturation.

- **Show available space**: makes it possible to show or hide available space on the chart and change the scale accordingly. Enable this option when the volume of the database becomes significant in relation to the available space.
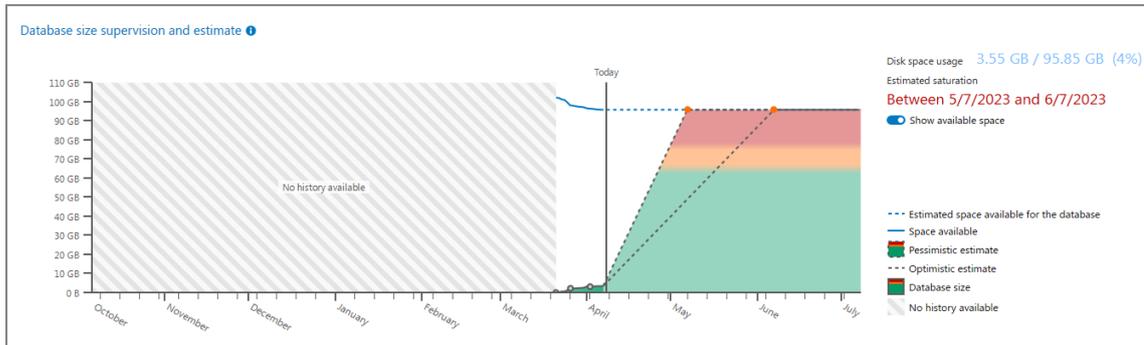
Managing log volume through the automatic degraded mode

To prevent the saturation of the log database when large quantities of logs are generated very quickly, a degraded mode is activated when the database reaches 81% of its maximum size. A red warning banner appears in the lower part of the administration console.

New agent and system logs sent to the Backoffice are no longer stored in the log database, but are permanently deleted.

However, if you have configured Syslog servers for agent managers, they will continue to receive agent logs.

To disable degraded mode and store logs again, reduce the log database volume until it is below the 81% threshold. To do so, follow the instructions below:

1. **Analyze logs** to understand the source of the logs,
2. **Adjust your security policy** to generate fewer logs, by reducing false positives, for example,
3. **Manually delete agent logs**,
4. Once you have reduced the log volume, click on **Back to standard mode** in the red banner of the administration console.
   The banner will disappear and logs will be stored in the log database again.

15.2.3 Configuring daily maintenance tasks

SES Evolution automatically conducts maintenance daily on the database to:

- Defragment indexes and optimize database performance,
- Delete the oldest logs in the log database. To set the criteria for deleting logs, refer to Managing the deletion of logs.

This maintenance task is scheduled by default to run at 00:00 at the backend's local time, but you can change the time. You are advised to schedule this maintenance before other SQL Server maintenance operations. This makes it possible to proceed with smaller backups later.

To configure daily maintenance tasks:

1. Select the **Backoffice > System** menu and click on the desired database tab, Administration or Logs.
2. In the upper banner, click on **Edit**.
3. In **Daily database maintenance**, enter the time at which you wish to run this maintenance task.
4. Click on **Save** at the top right of the window to save changes.

15.2.4 Managing the deletion of logs

Depending on the configuration of your security policies and the desired amount of logs, the log database can grow very quickly.

The following table estimates the average, minimum and maximum size of a log or a context:

	Average size	Minimum size	Maximum size
Log	4.5 KB	0.8 KB	10 KB
Simple context	100 KB	20 KB	600 KB
Full context	220 KB	20 KB	1200 KB

These figures are given for information. They can vary according to pool activity (e.g. applications used) and the security policies applied.

For more information on the on the contexts, Understanding what makes up a context.

There are various tools allowing you to delete logs and prevent the database from being saturated:

- Monthly deletion of system and agent logs that have exceeded the specified retention duration.
- Daily deletion of agent logs based on the specified criteria, performed by the maintenance task. For further information, refer to the section Configuring daily maintenance tasks.
- Manual deletion of agent logs. You can run a manual deletion task from the **System > Log Database panel** as described below, or from the **Agent Logs** panel as described in the Deleting events section.

Deleting system and agent logs monthly

System and agent logs are kept by default for 12 months, or 2 months if you use SQL Server Express. The oldest logs will be deleted. These values can be changed.

1. Select the **Backoffice > System** menu and click on the **Log database** tab.
2. In the upper banner, click on **Edit**.
3. In the section **Monthly deletion of logs**, choose how long (in months) you wish to keep agent events and system logs. After this duration expires, they will be deleted by an automatic task once a month.

To keep logs forever, disable monthly deletion by switching off the [⬤] button. Log deletion cannot be disabled if you use SQL Server Express.

Deleting agent logs daily

You can configure rules to delete logs daily. Every day, the maintenance task will check the criteria specified in the rules and delete logs that match all criteria:

To configure a rule to delete logs:

1. Select the **Backoffice > System** menu and click on the **Log database** tab.
2. In the upper banner, click on **Edit**.
3. In **Daily deletion of agent logs**, click on **Add a rule**.
   The **Add Daily Rule** wizard appears.
4. In the **General Filtering** panel, specify the criteria for deleting logs:

Date	You can choose whether to delete logs before a specific date, or delete all logs regardless of date.
Severity	You can select the log levels to be deleted. Refer to the section Managing logs.
Status	Select the status of the logs to delete. Refer to the section Managing logs.
Application	Choose whether the logs of all applications are affected, or indicate specific source or target applications.

5. Click **Next**.
6. In the **Event Type** panel, select the **Categories** or **Event Types** you want to delete, and click **Next**.
7. In the **Agents** panel, select the agent groups for which you wish to delete logs. By default, all agent groups are concerned.
8. Click on **OK**. The log deletion rule appears in the table.
9. If necessary, disable one or several rules by unselecting the **Status** checkbox to the left of the table. Click on **Hide disabled rules** to remove them from the table. The rules will be executed every day during the maintenance task. The results of the latest execution appear in the **Last result** column.

Deleting agent logs manually

If necessary, you can occasionally delete agent logs manually based on certain criteria.

1. Select the **Backoffice > System** menu and click on the **Log database** tab.
2. In the upper banner, click on **Edit**.
3. In **Manual deletion of agent logs**, click on **Configure and start manual deletion**.
   The **Add manual rule** wizard appears.

4. In the **General Filtering** panel, specify the criteria for deleting logs:

Date	Choose the logs you wish to delete by age (in days/months) or by date.
Severity	You can select the log levels to be deleted. Refer to the section **Managing logs**.
Status	Select the status of the logs to delete. Refer to the section **Managing logs**.
Application	Choose whether the logs of all applications are affected, or indicate specific source or target applications.

5. Click on **Next**.

6. In the **Event Type** panel, select the **Categories** or **Event Types** you want to delete, and click **Next**.

7. In the **Agents** panel, select the agent groups and then the agents for which you want to delete logs. By default, all agent groups and agents are concerned.

8. Click on **Estimate volume** at the bottom on the right to show the volume of logs that will be deleted.

9. Click on **Start**.
    All logs that meet the specified criteria will be deleted. The results of the latest execution appear in the **Last result** column.

> **ⓘ NOTE**
> The space freed up after the manual deletion will be shown in the **Database size supervision and estimate** chart only after the next maintenance task.

10. Click on **Run manual deletion again** to delete logs based on the same criteria used in the previous delete operation.

15.3 Configuring the Stormshield update server

Stormshield regularly provides several resources such as new security policies, rule sets, and Yara and IoC resources.

Updates are available by default on a Stormshield public server, but you can use a server of your choice.

To customize the update server:

1. Select the **Backoffice > System** menu, **General** tab.

2. Click on **Edit** in the upper banner.

3. The feature is enabled by default. If you disable it, the panel in which updates are downloaded will no longer be visible and notifications of new available updates will also be automatically disabled.

4. In the **Update server** section, select the frequency of connections to the server. Choose **Never** to disable automatic connection to the server. In this case, click on the **Check updates** button in the panel where updates are downloaded, to manually connect to the server.

5. Keep the Stormshield server's default address, or enter the address of a local server of your choice.

6. Ensure that the connection to the server functions by clicking on **Check address**.

> **ⓘ NOTE**
> The backend component sets up connections to the update server in HTTPS over TCP port 443. If there are several backend components, one of them will be chosen at random to set up the connection to the update server.

Once the server is configured, you can Downloading Stormshield updates.

15.4 Sending system log alerts by email

You can configure SES Evolution to send email alerts to recipients of your choice. Alerts are triggered by certain logs generated on SES Evolution backoffice components.

You must first configure an SMTP server. For more information, refer to Configuring an SMTP server.

You must hold the **Email Notifications-Modify** permission to configure the sending of alerts.

To send e-mail alerts:

1. In the **Backoffice > System** menu in the administration console, go to the **Email Notifications** tab.
2. Click on **Edit** in the upper banner.
3. In the **System log alerts** area, click on **Add rule**.
   The rule creation wizard opens.
4. Enter the rule settings:
   - **Rule name**.
   - **E-mail subject prefix** that the recipient will receive. By default, the subject of the e-mail begins with *SES EVOLUTION*. The prefix allows you to apply a specific process to SES Evolution alert e-mails in your mailbox.
   - **Log types** for which you want to trigger alerts from the **SES services**, **Databases**, and **External services** categories.
5. Click on **Next**.
6. In the field at the bottom of the screen, enter the e-mail address of the user who will receive the alerts, select the language, then click on **Add**.
7. Add more e-mail addresses if you wish to send alerts to several recipients.
8. Click on **Create**.
   The rule is added to the table in the **System log alerts** area.
9. Add other rules if necessary.

When system components generate logs corresponding to the rules, an email alert is sent.

You can disable or enable rules again by clicking on the checkbox in the **Enabled** column. The action buttons to the right of a rule can be used to duplicate or delete it.

You can temporarily stop emails being sent by disabling the **Enable notifications** option.

SES Evolution can also be used to send email alerts for agent logs or all dashboard content. For more information, see Sending agent logs alerts by email and Sending dashboard indicators by email.

15.5 Configuring an SMTP server

An SMTP server must be configured so that SES Evolution can send:

- **E-mail alerts based on the types of agent logs generated**,
- **Alert emails according to the type of system logs generated**,
- **Emails containing dashboard indicators**.

You must hold the **System-Modify** privilege to be able to configure an SMTP server.

1. In the **Backoffice > System** menu in the administration console, go to the **General** tab.
2. Click on **Edit** in the upper banner.
3. Define the following parameters in the **SMTP server** section:
   - **Server address**: Enter the DNS name or IP address of the SMTP server.
   - **Connection security:** Choose whether the connection needs to be encrypted, and by which protocol.
   - **Port**: Enter the communication port number, 587 by default for STARTTLS.
   - **Sender name**: Name of the notification sender. E-mail recipients will be able to see this name.
   - **Sender's e-mail address:** E-mail address of the notification sender. E-mail recipients will be able to see this address. Be sure to check the sender's mailbox regularly to see if any emails are being returned due to errors in the recipient's address.
   - **Authentication required**: Enable this option if your SMTP service requires authentication. If so, enter the ID and password of the service.
4. Click on **Check settings** to test the connection to the SMTP server. No e-mails will be sent during this test.

16. Enabling and managing SES Evolution's public API

SES Evolution has a standard REST API with which the solution can be used through your own orchestration tools.

Some SES Evolution features are not yet available in the public API. It will be enriched along with every new version.

The public API is not enabled by default.

Authentication over the public API is secured by API keys that administrators generate. The usage and validity of these keys can be configured.

POST operations performed via the public API are recorded in system logs.

To facilitate the use of the API, you can access OpenAPI documentation through a link shown in the administration console. It can also be found on the Stormshield Technical Documentation website.

16.1 Requirements

You need to hold the **Public API-Modify** permission to enable the public API and generate keys from the **API keys** menu in the administration console. See Managing users on the SES Evolution administration console.

This permission has priority over the **System** permission. If an administrator does not hold any **System** permissions, but holds the **Public API-Show** or **Public API-Modify** permission, the **System** menu will appear with only the **API keys** tab.

16.2 Enabling the public API

SES Evolution's public API is disabled by default.

When the public API is enabled:

- Access to routes on the API is allowed,
- Authorized administrators can create, modify and revoke API keys.

To enable the public API:

1. In **Backoffice** > **system**, show the **API keys** tab.
2. Click on **Edit** in the upper banner.

3. Select **Enable Public API**.
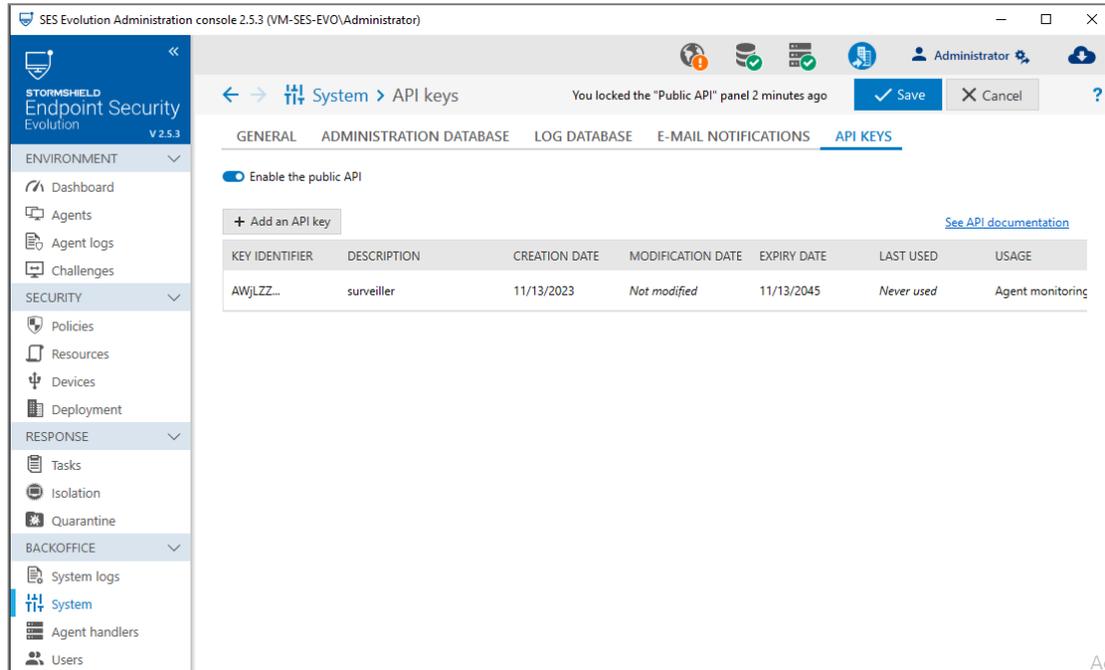
16.3 Adding an API key

The *API keys* tab makes it possible to add, modify and revoke keys that grant access to SES Evolution's public API routes.

These keys have an ID, description, creation date, expiry date and usage. They are required for every public API request.

To add a key:

1. Click on **Edit** in the upper banner.
2. Click on **Add an API key**.
3. Enter a description.
4. Select a validity duration. Once the key has been created, you will no longer be able to modify this duration.
5. Check the usage of the key, which open access to the various routes offered by the API.
6. When you click on **OK**, the key will not be saved in the database. It must be copied and stored in a safe place as it will no longer be available later.
7. Click on **Close**, then on **Save** in the upper banner.

All SES Evolution administrators who are allowed to display the **API keys** tab have access to the list of keys created.

API keys can be used in API documentation to test requests:

1. Click on the **See the API documentation** link.

2. Click the browse button [Authorize 🔓].
3. Enter the API key in the **Value** field.
4. Click on **Authorize** then on **Close**.

16.4 Revoking an API key

API keys can be revoked. Once they are revoked, you can no longer use them to submit requests on the API.

To revoke a key:

1. Click on **Edit** in the upper banner.

2. On the line of the key to revoke, click on the icon ☒ in the **Actions** column.

3. Confirm.

If you wish to display only unrevoked keys in the table, enable the **Hide revoked keys** button.

16.5 Troubleshooting

16.5.1 The public API documentation does not appear

Situation: The **See API documentation** link in the administration console opens a web browser and the documentation does not appear.

Cause: If you have upgraded SES Evolution to version 2.7.2 and your backend server is installed on the Windows Server 2022 operating system, the TLS 1.3 option is enabled by default in the backend IIS settings. This option makes the API documentation incompatible with this operating system.

Solution: In your IIS service manager, disable the **TLS 1.3 over TCP** option in the settings of the backend host name. The host name to be modified can be seen in the URL of the API documentation, which corresponds to the host name of the backend used when SES Evolution was installed.

17. Troubleshooting

If you encounter issues with backoffice components or with agents on user workstations, SES Evolution allows you to troubleshoot with a dedicated tool.

If you encounter issues only with agents, SES Evolution also offers a challenge mechanism that allows you to disable or uninstall agents.

17.1 Resolving issues with challenges

When users encounter issues on their workstations or need to perform operations that they cannot perform while the SES Evolution agent is running, they can ask the security administrator to temporarily disable or uninstall the agent, or allow it to run a diagnostic.

The **Challenges** feature must be enabled in the **Settings** tab of agent groups:

> ✍ **EXAMPLES**
>
> - Self-protection may need to be disabled on the agent in order to debug potential compatibility issues with other programs.
> - The agent may need to be temporarily stopped on an offline workstation for the duration of maintenance operations such as the installation or update of an ERP.

As the security administrator, your responsibility is to choose which action to run on the user's workstation and you must hold a role that includes the **Challenge-response** privilege.

The challenge mechanism is based on a question/answer system between the agent and the console.

A user on the workstation generates a character string (the question) from the agent that they communicate to you by telephone or email. You then enter this string in the console, which generates another character string (the response) containing the definition of the action to allow. You forward this response to the user so they can enter it in the agent's interface. The action will then be allowed for the duration that you have defined.

The mechanism functions even when the agent is not connected to the network.

Three operations are possible using challenges:

- Enabling maintenance mode,
- Stopping the agent,
- Running the diagnostic with or without tracing,
- Uninstalling the agent.

Administration privileges are not required on the user's workstation to enable these operations via the challenge mechanism.

For further information on Maintenance mode, refer to the section **Understanding self-protection on agents and performing maintenance operations**.

For more information on diagnostics and tracing, refer to the section **Diagnosing issues on agents**.

17.1.1 Enabling Maintenance mode

Maintenance mode disables self-protection on the agent so that maintenance operations or tests can be conducted.

> ✎ **EXAMPLE**
> You can use this mode to change permissions on certain registry keys.

To enable this mode using a challenge, ask the user to:

1. Open the agent interface by double-clicking on 🔲 in the taskbar.
2. Go to 🧰 to open the **Help and support** panel.
3. Click on **Request a challenge** in the **Helpdesk** tab.
4. Send you the challenge code generated.
5. Keep the **New challenge** window open.

On your side:

1. Open the **Environment > Challenges** menu in the console.
2. Enter the challenge code.
3. Select **Maintenance mode**.
4. Select a duration.
5. Click on **Generate**.
6. Send the response code to the user.
7. Ask the user to enter the response code in the **New challenge** window, then click on **Start the challenge**.

The user can stop the challenge in progress at any time in the lower banner in the agent's interface.

Maintenance mode can also be enabled in the **Helpdesk** tab in **Help and support** panel of the agent's interface. It must be enabled beforehand in the agent group configuration, and the user of the workstation must have administration privileges to enable this mode. For further information on disabling self-protection and Maintenance mode, refer to the section Understanding self-protection on agents and performing maintenance operations.

17.1.2 Stopping an agent

If some issues on an agent persist even though you enabled Maintenance mode, the agent may need to be stopped temporarily. Stopping the agent makes it possible to disable the protection applied by rules in the security policy in addition to self-protection rules.

To stop an agent using a challenge, ask the user to:

1. Open the agent interface by double-clicking on 🔲 in the taskbar.
2. Go to 🧰 to open the **Help and support** panel.
3. Click on **Request a challenge** in the **Helpdesk** tab.
4. Send you the challenge code generated.
5. Keep the **New challenge** window open.

On your side:

1. Open the **Environment > Challenges** menu in the console.
2. Enter the challenge code.
3. Select **Stop agent**.
4. Select a duration.
5. Click on **Generate**.
6. Send the response code to the user.
7. Ask the user to enter the response code in the **New challenge** window, then click on **Start the challenge**.

The user can stop the challenge in progress at any time in the lower banner in the agent's interface.

17.1.3 Running a diagnostic

SES Evolution provides a diagnostic tool with which data on user workstations can be collected when an issue occurs. Stormshield's technical support team can then analyze this data.

To run a diagnostic from an agent using a challenge, ask the user to:

1. Open the agent interface by double-clicking on in the taskbar.
2. Go to to open the **Help and support** panel.
3. Click on **Request a challenge** in the **Helpdesk** tab.
4. Send you the challenge code generated.
5. Keep the **New challenge** window open.

On your side:

1. Open the **Environment > Challenges** menu in the console.
2. Enter the challenge code.
3. Select **Run diagnostic**.
4. Enable tracing if necessary, as it will allow you to record the series of operations the user performed that led to the unexpected behavior on SES Evolution.
5. Click on **Generate**.
6. Send the response code to the user.
7. Ask the user to enter the response code in the **New challenge** window, then click on **Start the challenge**.
8. Next, refer to Diagnosing issues on agents to use the diagnostic tool.

17.1.4 Uninstalling an agent

If another program is incompatible with the SES Evolution agent and prevents the user from working, for example, the only solution is to uninstall the agent.

To uninstall an agent using a challenge, ask the user to:

1. Open the agent interface by double-clicking on in the taskbar.
2. Go to to open the **Help and support** panel.
3. Click on **Request a challenge** in the **Helpdesk** tab.

4. Send you the challenge code generated.

5. Keep the **New challenge** window open.

On your side:

1. Open the **Environment > Challenges** menu in the console.

2. Enter the challenge code.

3. Select **Uninstall the agent**.

4. Click on **Generate**.

5. Send the response code to the user.

6. Ask the user to enter the response code in the **New challenge** window, then click on **Start the challenge**.

Once the challenge has started, it cannot be stopped and there is no way to go back. Users must restart their workstations to properly end the process.

The agent can also be uninstalled through the standard procedure of uninstalling programs, provided that the agent group conf allows it. Administration privileges are required. For further information, refer to the sections Allowing administrators to uninstall agents and Uninstalling agents.

17.2 Troubleshooting issues

When backoffice components (backend server, agent handler and administration console) or SES Evolution agents do not function normally, Stormshield Technical support may suggest that you use the troubleshooting tool provided with the solution. This tool collects data concerning the component causing issues and the workstation's system. Technical support has a tool to analyze this data, which is compiled in the form of a diagnostic package, and can therefore locate the cause of the issue.

17.2.1 Diagnosing issues on backoffice components

The *EsDiag.exe* diagnostic tool is installed with all SES Evolution backoffice components. It allows you to collect varied data about the workstation on which an issue occurred and for the data to be gathered in a diagnostic package (*.zip*) to be provided to Stormshield Technical Support for analysis.
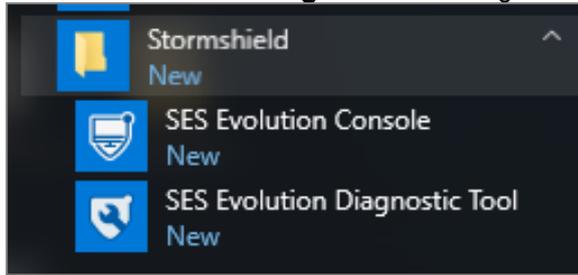
> **ℹ️ NOTE**
> Administration privileges are required to use the diagnostic tool.
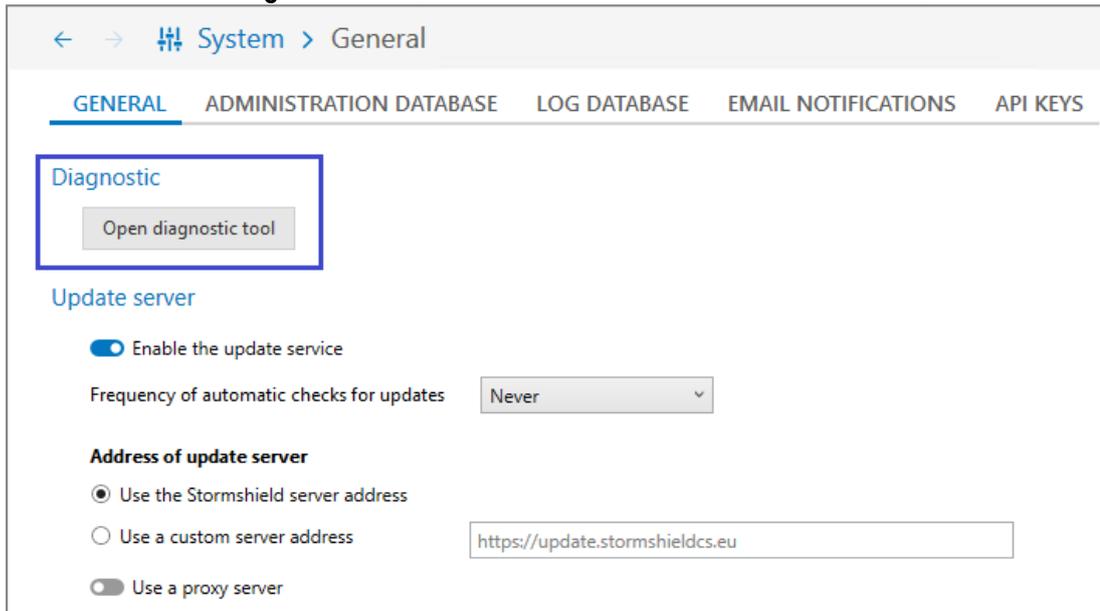
Opening the diagnostic tool

Depending on the SES Evolution components installed on the workstation to be analyzed, there are four ways to open the diagnostic tool:

- Via the **SES Evolution Diagnostic Tool** entry in the Windows **Start** menu,

,

- Via the *EsDiag.exe* executable file in the installation folder of a component,
- Via the **Backoffice > System** menu, **General** tab in the administration console,

- Via command line options.

You can start only one data collection at a time per workstation.

Using the diagnostic tool in the graphical interface

You have opened the diagnostic tool in one of the first three ways suggested above. To use it:

1. On the home screen, select the checkbox to accept the collection of personal data.

2. If an agent in version 2.4 or higher is installed on the workstation, the following screen will offer to diagnose the agent and the backoffice, or only the backoffice. The first option offers a more comprehensive collection scope. If you choose the first option, click on **Troubleshoot from the agent** to open the agent's interface. The Collecting diagnostic data option must have been enabled beforehand in the agent group, otherwise you must use the challenge mechanism to run the diagnostic from the agent. To proceed, refer to the section Diagnosing issues on agents.

3. If you have chosen to collect data only from the backoffice, on the next screen, indicate the destination folder and name of the diagnostic package. Add a description if necessary.

4. Click on **Start**. The tool collects data about the backoffice components and the workstation's system.

5. During data collection or compression, if you cancel the process, close the diagnostic tool or log out of Windows, any data already collected will be deleted.

6. At the end of the collection, you can check **Open file location when closing the tool** and close.

Using the diagnostic tool in the command line interface

You can also use the following commands to use the diagnostic tool:

`EsDiag /GenerateDiagnostic <path.zip> /AcknowledgePersonalDataCollection /DiagnosticComment <comment>`	Generates a diagnostic package. Specify the destination folder of the package. The `/AcknowledgePersonalDataCollection` parameter is mandatory to allow the collection of personal data. The `/DiagnosticComment` parameter is optional, and makes it possible to add comments.
`EsDiag /CancelDiagnostic`	Cancels the diagnostic in progress.

17.2.2 Diagnosing issues on agents

The diagnostic tool can be started from the agent interface on the users' workstations.

It allows you to collect varied data about the workstation on which an issue occurred and for the data to be gathered in a diagnostic package (*.zip*) to be provided to Stormshield Technical Support for analysis.

In addition to the diagnostic data, you can request tracing to record the series of operations the user performed that led to the unexpected behavior on SES Evolution.

Tracing collects debugging traces from each module of the agent.

You can start only one data collection at a time per workstation.

Using the diagnostic tool in the graphical interface

To use the diagnostic tool on a user workstation, ask the user to perform the following operations:

1. Double-click on the [icon] icon in the task bar to open the agent's interface.
2. Click on the [icon] tab to open the **Help and support** panel.
3. Click on the **Diagnosis** tab.
4. By default, you must use the challenge mechanism to allow the user to run the diagnostic tool. Ask the user to generate a challenge code in the **Helpdesk** tab in the **Help and support** panel. For further information on challenges, refer to the section **Resolving issues with challenges**. The **Challenges** feature must be enabled in the **Settings** tab of agent groups: The contents of the **Diagnosis** tab may vary depending on the settings of the agent group and whether the user is an administrator. To find out more, refer to the information at the end of this procedure.

   > **🛈 NOTE**
   > Administration privileges are not required to start a diagnostic via a challenge.

5. If you have allowed tracing in the challenge, ask the user to start by clicking **Start tracing**.
6. Ask the user to reproduce the scenario that caused the abnormal behavior and then click **Next**.

7. The user must then specify the destination folder and the name of the diagnostic package. The user can enter a description if necessary.

8. The user clicks on **Next**. The tool collects data about the SES Evolution agent and the workstation's system. During data collection or compression, if the user cancels, closes the diagnostic tool or logs out of Windows, any data already collected will be deleted.

9. At the end of the collection, the user can check **Open file location when closing the tool** and close.

10. Ask the user to provide the generated *.zip* diagnostic package for Stormshield Technical Support to analyze.

For users who have administration privileges on their workstations, the **Allow collection of diagnostic data** option in the agent group **Settings** tab allows them to start a diagnostic directly from the agent, without going through a challenge. For more information, see the section Collecting diagnostic data.

When the challenge feature is disabled in the agent group **Settings** tab and data collection is allowed in agent groups as well, non-admin users of their workstation can request privilege escalation to start a diagnostic from the agent.

Using the diagnostic tool in the command line interface

You can also establish a diagnosis on a user's workstation via a script, by launching EsGui ([…]\\Stormshield\SES Evolution\Agent\Bin\Gui) with the following command line options. Administration privileges are required and the Collecting diagnostic data option must be enabled in the agent group.

`EsGui /GenerateDiagnostic`
`<path.zip>`
`/AcknowledgePersonalDataCollection`
`/DiagnosticComment <comment>`
	Generates a diagnostic package without tracing. Specify the destination folder of the package. The `/AcknowledgePersonalDataCollection` parameter is mandatory to allow the collection of personal data. The `/DiagnosticComment` parameter is optional, and makes it possible to add comments.

`EsGui /StartDiagnosticWithTraces` `/AcknowledgePersonalDataCollection`	Starts diagnosing and tracing.

`EsGui /StopDiagnosticWithTraces`
`<path.zip> /DiagnosticComment`
`<comment>`
	Stops tracing and finishes generating the diagnostic package. Specify the location to save the package.

`EsGui /CancelDiagnostic`	Cancels the diagnostic in progress.

For more information on EsGui, refer to the section Using the EsGui command.

18. Further reading

Additional information and answers to questions you may have about SES Evolution are available on the Documentation website and in the Stormshield knowledge base (authentication required).

Appendix A. Supported OSSEC functions

There are several differences between the SES Evolution analysis engine and OSSEC, especially with regard to supported configuration items. This appendix indicates whether SES Evolution supports each listed OSSEC decoder or rule item.

For more information on using the analysis engine, refer to Importing OSSEC security rules.

A.1 Decoder file items

Supported decoder items

Configuration item	Remarks
`<decoder name="...">`	The decoder name is mandatory.
`<decoder name="...">` `<parent>...</parent>` `</decoder>`	Makes it possible to link the decoder to a higher level decoder. **ⓘ NOTE** SES Evolution allows more than two levels of decoders.
`<decoder name="...">` `<prematch>...</prematch>` `</decoder>`	Advanced OSSEC regular expression that can be used to quickly verify whether the decoder is suitable for the log message.
`<decoder name="...">` `<prematch_pcre2>...</prematch_pcre2>` `</decoder>`	PCRE2 regular expression that can be used to quickly verify whether the decoder is suitable for the log message.
`<decoder name="...">` `<program_name>...</program_name>` `</decoder>`	Simple OSSEC regular expression targeting the *program_name* field extracted during the pre-decoding phase, which can be used to quickly verify whether the decoder is suitable for the log message.
`<decoder name="...">` `<program_name_pcre2>...</program_name_pcre2>` `</decoder>`	PCRE2 regular expression targeting the *program_name* field extracted during the pre-decoding phase, which can be used to quickly verify whether the decoder is suitable for the log message.
`<decoder name="...">` `<regex>...</regex>` `<order>...</order>` `</decoder>`	Extracts fields from the log using an advanced OSSEC regular expression with capture groups. SES Evolution makes it possible to extract to any field name.
`<decoder name="...">` `<pcre2>...</pcre2>` `<order>...</order>` `</decoder>`	Extracts fields from the log using a PCRE2 regular expression with capture groups. SES Evolution makes it possible to extract to any field name.
`<decoder name="...">` `<use_own_name>...</use_own_name>` `</decoder>`	Makes it possible to write rules later that target the name of this decoder when it is not at the first level. SES Evolution ignores this option but supports decoders from all levels in the *decoded_as* option in rules.

`<decoder name="...">` `<type>...</type>` `</decoder>`	Allows the decoder to be classified. The supported values are: *firewall, ids, web-log, syslog, squid, windows, host-information* and *OSSEC*. The first seven mandatory rules (in *rules_config.xml*) correspond to all of these types except *host-information*.
`<decoder name="...">` `<fts>...</fts>` `</decoder>`	Allows *n-tuple* fields to be cached to see whether their values have already been observed together.

Unsupported decoder items

Configuration item	Remarks
`<decoder status="...">`	Even though OSSEC contains code to read this field, any configuration that contains it is invalid.
`<decoder id="...">`	OSSEC contains code to read this field, but does not use its value.
`<decoder type="...">`	OSSEC contains code to read this field, but does not use its value.
`<decoder name="...">` `<plugin_decoder>...</plugin_decoder>` `</decoder>`	Allows users to compile their own decoders for specific needs.
`<decoder name="...">` `<accumulate/>` `</decoder>`	Supports logs that span several lines with common fields.

A.2 Rule file items

Supported rule items

Configuration item	Remarks
`<var name="FREQ">8</var>` … `<group name="...">` `<rule ... frequency="$FREQ">` …	Declares constants at the top of the file.
`<group name="...">` `<rule ...>` … `</rule>` `</group>`	`<rule>` items must be under a `<group>` item. The *name* attribute is mandatory and ends with a comma. Used in classifying rules found in the group.
`<rule id="123456">`	The *id* attribute of a rule is mandatory, with a value between 1 and 999999 inclusive.
`<rule overwrite="yes|no">`	Makes it possible to do away with a unique *id* attribute, replaces a rule that was defined earlier.
`<rule level="0..15">`	Mandatory. Assigns a level of severity to the rule; level 0 rules are evaluated on a higher priority than others.

`<rule accuracy="0">`	Gives rules containing this attribute lower priority than others.
`<rule maxsize="0..9999">`	Makes it possible for the rule to apply only to logs with a message that is at least as long as the value of this attribute.
`<rule timeframe="..." frequency="...">`	Declares a composite rule that is triggered if an event occurs several times within the defined time frame.
`<rule noalert="...">`	Considers that a rule does not apply if no child rules apply.
`<rule ignore="...">`	Inhibits the rule for a set number of seconds after it is triggered.
`<rule id="..." level="..."> <decoded_as>...</decoded_as> </rule>`	Indicates the first-level decoder (or second-level using the *use_own_name* option) that must have been used for the message. SES Evolution supports decoder names from all levels and ignores the *use_own_name* option.
`<rule id="..." level="..."> <if_sid>...</if_sid> </rule>`	Links a rule to a parent rule with a rule ID.
`<rule id="..." level="..."> <if_group>...</if_group> </rule>`	Links a rule to parent rules with a group name.
`<rule id="..." level="..."> <if_level>...</if_level> </rule>`	Links a rule to parent rules with a minimum level of severity.
`<rule id="..." level="..."> <regex>...</regex> </rule> <rule id="..." level="..."> <match>...</match> </rule> <rule id="..." level="..."> <pcre2>...</pcre2> </rule> <rule id="..." level="..."> <match_pcre2>...</match_pcre2> </rule>`	Simple/advanced OSSEC/PCRE2/PCRE2 regular expression targeting the log message to determine whether the rule matches. **NOTE** The last two variants are synonymous.
`<rule id="..." level="..."> <user>...</user> </rule> <rule id="..." level="..."> <user_pcre2>...</user_pcre2> </rule>`	Simple OSSEC/PCRE2 regular expression targeting the *srcuser* decoded field, or if there isn't one, the *dstuser* decoded field, to determine whether the rule matches.

`<rule id="..." level="...">`
`<srcip>...</srcip>`
`</rule>`
`<rule id="..." level="...">`
`<dstip>...</dstip>`
`</rule>`
	IPv4 or IPv6 address specification (individual addresses, ranges, networks with mask length, etc.) compared to the *srcip* or *dstip* fields to determine whether the rule matches.
	The specification can be expressed as a negative by placing an exclamation mark in front of it.

`<rule id="..." level="...">`
`<srcport>...</srcport>`
`</rule>`
`<rule id="..." level="...">`
`<srcport_pcre2>...</srcport_pcre2>`
`</rule>`
`<rule id="..." level="...">`
`<dstport>...</dstport>`
`</rule>`
`<rule id="..." level="...">`
`<dstport_pcre2>...</dstport_pcre2>`
`</rule>`
	Simple OSSEC/PCRE2 regular expressions targeting the *srcport* and *dstport* decoded fields to determine whether the rule matches.

`<rule id="..." level="...">`
`<id>...</id>`
`</rule>`
`<rule id="..." level="...">`
`<id_pcre2>...</id_pcre2>`
`</rule>`
	Simple OSSECPCRE2 regular expression targeting the *id* decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
`<status>...</status>`
`</rule>`
`<rule id="..." level="...">`
`<status_pcre2>...</status_pcre2>`
`</rule>`
	Simple OSSEC/PCRE2 regular expression targeting the *status* decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
`<hostname>...</hostname>`
`</rule>`
`<rule id="..." level="...">`
`<hostname_pcre2>...</hostname_pcre2>`
`</rule>`
	Simple OSSEC/PCRE2 regular expression targeting the *hostname* pre-decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
`<extra_data>...</extra_data>`
`</rule>`
`<rule id="..." level="...">`
`<extra_data_pcre2>...</extra_data_pcre2>`
`</rule>`
	Simple OSSEC/PCRE2 regular expression targeting *data* decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
 `<program_name>...</program_name>`
 `</rule>`
 `<rule id="..." level="...">`
 `<program_name_pcre2>...</program_name_pcre2>`
 `</rule>`
	Simple OSSEC/PCRE2 regular expression targeting the *program_name* pre-decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
 `<url>...</url>`
 `</rule>`
 `<rule id="..." level="...">`
 `<url_pcre2>...</url_pcre2>`
 `</rule>`
	Simple OSSEC/PCRE2 regular expression targeting the *url* decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
 `<action>...</action>`
 `</rule>`
	Exact value compared to the *action* decoded field to determine whether the rule matches.

`<rule id="..." level="...">`
 `<field name="...">...</field>`
 `</rule>`
	Advanced OSSEC regular expression targeting the decoded field indicated, to determine whether the rule matches.

`<rule id="..." level="...">`
 `<time>...</time>`
 `</rule>`
	Specifies the time slot during which the rule applies. Supports any format that OSSEC supports.
	
	 ✏️ **EXAMPLE**
	 `<time>1:30-17:45</time>` ; `<time>1 am - 12:30 PM</time>` ; `<time>!08:00-17:30</time>`
	
	 SES Evolution uses the system time zone to evaluate local time.

`<rule id="..." level="...">`
 `<weekday>...</weekday>`
 `</rule>`
	Specifies the days of the week when the rule is enabled. Supports any format that OSSEC supports.
	
	 ✏️ **EXAMPLE**
	 `<weekday>wed fri sun</weekday>` ; `<weekday>weekdays sunday</weekday>` ; `<weekday>! tue wed</weekday>`
	
	 SES Evolution uses the system time zone to evaluate local time, and therefore the day.

`<rule id="..." level="...">`
 `<cve>...</cve>`
 `</rule>`
 `<rule id="..." level="...">`
 `<info type="cve">...</info>`
 `</rule>`
	Describes the rule by associating it with a known vulnerability.

`<rule id="..." level="...">`
`<info type="text">...</info>`
`</rule>`
`<rule id="..." level="...">`
`<info type="link">...</info>`
`</rule>`
`<rule id="..." level="...">`
`<info type="osvdb">...</info>`
`</rule>`
	Describes the rule using text, a link or an Open Source Vulnerability Database item.
	SES Evolution supports only a single item of each type in the same rule.

`<rule id="..." level="...">`
`<group>...</group>`
`</rule>`
	Adds groups that the rule belongs to, in addition to those specified in the rule's parent `<group>` node.

`<rule id="..." level="...">`
`<description>...</description>`
`</rule>`
	Mandatory description of the event to which the rule applies. SES Evolution uses this description in the log summary that appears in the agent and console.

`<rule id="..." level="...">`
`<category>...</category>`
`</rule>`
	Links the rule to one of the decoder types. Option used for rules 1 to 7 in the file *rules_config.xml*.

`<rule id="..." level="...">`
`<if_fts/>`
`</rule>`
	Makes the rule effective only if a decoder has detected (with the fts option) that a set of fields had values seen together for the first time.

`<rule id="..." level="...">`
`<ignore>...</ignore>`
`</rule>`
`<rule id="..." level="...">`
`<check_if_ignored>...</check_if_ignored>`
`</rule>`
	Makes it possible to cache sets of field values, and disable a rule later for the same value sets.

`<rule id="..." level="...">`
`<check_diff/>`
`</rule>`
	Allows you to ignore two consecutive and identical logs.

`<rule id="..." level="..." frequency="..." timeframe="...">`
`<if_matched_regex>...</if_matched_regex>`
`</rule>`
	Makes it possible to trigger a composite rule if several logs generated recently can be described with an advanced OSSEC regular expression.

`<rule id="..." level="..." frequency="..." timeframe="...">`
`<if_matched_group>...</if_matched_group>`
`</rule>`
	Makes it possible to trigger a composite rule if several logs generated recently were described with a rule in a given group

`<rule id="..." level="..." frequency="..." timeframe="...">`
`<if_matched_sid>...</if_matched_sid>`
`</rule>`
	Makes it possible to trigger a composite rule if several logs generated recently were described with a rule that has a given ID.

`<rule id="..." level="..." frequency="..." timeframe="...">` `<same_source_ip/>` `</rule>` `<rule id="..." level="..." frequency="..." timeframe="...">` `<same_src_port/>` `</rule>` `<rule id="..." level="..." frequency="..." timeframe="...">` `<same_dst_port/>` `</rule>` `<rule id="..." level="..." frequency="..." timeframe="...">` `<same_id/>` `</rule>` `<rule id="..." level="..." frequency="..." timeframe="...">` `<same_user/>` `</rule>`	Makes it possible to trigger a composite rule if several logs sharing the same *srcip*, *srcport*, *dstport*, *id* or *user* field are found.
`<rule id="..." level="..." frequency="..." timeframe="...">` `<different_srcip/>` `</rule>` `<rule id="..." level="..." frequency="..." timeframe="...">` `<different_url/>` `</rule>`	Makes it possible to trigger a composite rule if several logs with separate vales are found for the same *srcip* or *url* field.

Unsupported rule items

Configuration item	Supported	Remarks
`<rule id="..." level="...">` `<srcgeoip>...</srcgeoip>` `</rule>` `<rule id="..." level="...">` `<srcgeoip_pcre2>...</srcgeoip_pcre2>` `</rule>` `<rule id="..." level="...">` `<dstgeoip>...</dstgeoip>` `</rule>` `<rule id="..." level="...">` `<dstgeoip_pcre2>...</dstgeoip_pcre2>` `</rule>`	No	SES Evolution does not use the *libgeoip library*.
`<rule id="..." level="...">` `<list lookup="..." field="...">...</list>` `</rule>`	No	Quick search for a field in a CDB (container database). SES Evolution does not use the CDB library.

<rule id="..." level="..."> <compiled_rule>...</compiled_rule> </rule>	Partial	Allows users to compile their own rules for specific needs. SES Evolution supports the *is_simple_http_request* function, which serves as an example, but is used in standard rule sets.
<rule id="..." level="..."> <options>...</options> </rule>	Partial	SES Evolution supports only the *no_log* option; e-mail alerts or active responses are not supported.
<rule id="..." level="..." frequency="..." timeframe="..."> <not_same_source_ip/> </rule> <rule id="..." level="..." frequency="..." timeframe="..."> <not_same_id/> </rule> <rule id="..." level="..." frequency="..." timeframe="..."> <not_same_user/> </rule>	No	Unnecessary OSSEC options in a configuration with the sole purpose of canceling a <*same_...*> option written earlier in the same rule: you are advised to remove the previous option.
<rule id="..." level="..." frequency="..." timeframe="..."> <same_location/> </rule> <rule id="..." level="..." frequency="..." timeframe="..."> <not_same_agent/> </rule>	No	SES Evolution analyzes and correlates logs at the agent level instead of the server level. As a result, multiple agent logs cannot be correlated; these options will therefore be ignored.
<rule id="..." level="..." frequency="..." timeframe="..."> <different_srcgeoip/> </rule>	No	SES Evolution does not use the *libgeoip library*.

STORMSHIELD

documentation@stormshield.eu

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

Copyright © Stormshield 2026. All rights reserved. All other company and product names contained in this document are trademarks or registered trademarks of their respective companies.