

RSA NetWitness

Version 11.7

Alerting with ESA Correlation Rules User Guide



Contact Information

NetWitness Community at <https://community.netwitness.com> contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA Conference Logo, RSA, and other trademarks, are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to <https://www.rsa.com/en-us/company/rsa-trademarks>. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates and are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

March 2022

Contents

Getting Started with ESA	11
How ESA Generates Alerts	12
Data Source Configuration Changes	12
An Endpoint Risk Scoring Rules Bundle is available in NetWitness Platform	12
Best Practices	13
Understand Event Stream Analysis Rule Types	13
Best Practices for Writing Rules	14
Best Practices for Working with RSA Live Rules	15
Best Practices for Deploying Rules	15
Best Practices for System Health	16
Troubleshoot ESA	17
Troubleshoot ESA Correlation Services	17
Troubleshoot RSA Live Rules for ESA	19
Troubleshoot ESA Rules	20
SMTP Notification Error Example	27
Integration-Server SMTP Notification Error Example	27
Example ESA Correlation Server Warning Message for Missing Meta Keys	27
Multi-Valued Warning Message Example	28
Single Value Warning Message Example	28
Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_ source_id	28
Steps to Troubleshoot Memory Issues with an ESA Service Offline	30
Step 1: Verify that your Host Is Running	30
Step 2: View Detailed Statistics in Health & Wellness	31
Step 3: Bring up your ESA Services	35
Step 4: Check the Alerts and Events Volume	35
View Alert Summaries	35
View Events Matched	36
Step 5: Disable and Repair the Rule that Caused Issues	37
Disable Rules	37
Edit Rules	37
Deploy Rules	37
Verify that the Rules are Enabled	37
(Optional) Check the ESA Correlation Log Files for More Information	38
ESA Rule Troubleshooting with Nw-Shell	38
Find Your Engine Name for Nw-Shell	38

Connect to an ESA Correlation Server	39
View the Contents of a Named Window	39
See the Method Input and Output	40
Obtain Correlation Server Metrics for ESA Rule Deployment Troubleshooting Using Nw-Shell	41
View Memory Metrics for Rules	44
Prerequisites	44
View Health Statistics and Trends for ESA Correlation in New Health & Wellness	44
View Memory Metrics for an ESA Correlation Service in Health & Wellness	45
View Memory Metrics for an ESA Correlation Service and its ESA Rules	47
How ESA Handles Sensitive Data	50
How ESA Treats Sensitive Data from Core Services	50
Advanced EPL Rule	50
Enrichment Source	51
How to Remove Sensitive Meta Keys Globally from All Alerts	51
ESA Rule Types	53
Sample Rules	53
Endpoint Risk Scoring Rules Bundle	53
Trial Rules Mode	54
ESA Permissions	54
Practice with Sample Rules	55
Rule Library	56
Practice with Sample Rules	57
Working with Trial Rules	58
Deploy Rules as Trial Rules	58
Add Rules to the Rule Library	60
Download Configurable RSA Live ESA Rules	61
Prerequisites	61
Download RSA Live ESA Rules	61
Customize an RSA Live ESA Rule	64
Prerequisites	65
Configure Parameters for an RSA Live ESA Rule	65
Add a Rule Builder Rule	66
Step 1. Name and Describe the Rule	66
Prerequisites	66
Name and Describe a Rule	66
Step 2. Build a Rule Statement	68
Example	68
Prerequisites	68
Build a Rule Statement	69

To Add a Whitelist	71
To Add a Blacklist	71
Example: Blacklist	72
Example: Strict Pattern Matching and Using the Is Not Null Operator	73
Example Results	76
Example: Grouping the Rule Results	77
Example: Working with Numeric Operators	78
Step 3. Add Conditions to a Rule Statement	80
Example	80
Add Conditions to a Rule Statement	80
Example	81
Validate an ESA Rule	82
Working with Rules	86
Edit, Duplicate or Delete a Rule	86
Edit a Rule	86
Duplicate a Rule	86
Delete a Rule	86
Filter or Search for Rules	87
Prerequisites	87
Filter Rules	87
Search for Rules	88
Import or Export Rules	88
Import ESA Rules	88
Export ESA Rules	89
Choose How to be Notified of Alerts	90
Notification Methods	90
Email Notifications	91
Syslog	91
Script Alerter	91
Add Notification Method to a Rule	92
Prerequisites	92
Add a Notification Method to a Rule	92
Add a Data Enrichment Source	94
Example Rule with Enrichments	94
Enrichment Sources	96
Configure a Context Hub List as an Enrichment Source	97
Prerequisites	97
Configure a Context Hub List as an Enrichment Source	97
Configure an In-Memory Table as an Enrichment Source	99
Prerequisites	100

Configure an Ad hoc In-Memory Table	100
Add a Recurring In-Memory Table	103
Add an Enrichment to a Rule	103
Deploy Rules to Run on ESA	105
How an ESA Rule Deployment Works	105
ESA Rule Deployment Steps	106
Step 1. Add an ESA Rule Deployment	106
Prerequisites	106
Step 2. Add an ESA Service	108
Step 3. Add Data Sources	110
Step 4. Add and Deploy Rules	115
(Optional) Add a Data Source Filter	117
Deploy the Endpoint Risk Scoring Rules Bundle	120
Additional ESA Rule Deployment Procedures	121
Replace an ESA Service in an ESA Rule Deployment	121
Remove an ESA Service from an ESA Rule Deployment	121
Add an ESA Service to an ESA Rule Deployment	122
Edit a Data Source in an ESA Rule Deployment	122
Add or Remove a Data Source	124
Remove a Data Source from an ESA Rule Deployment	124
Add a Data Source to an ESA Rule Deployment	124
(Optional) Add or Remove a Data Source Filter	125
Remove a Data Source Filter from an ESA Rule Deployment	125
Add a Data Source Filter to an ESA Rule Deployment	125
Adjust Maximum Sessions for the Data Source Filter	126
Edit or Delete a Rule in a Deployment	126
Edit a Rule	126
Delete a Rule	127
Stop Deployment	127
Edit the ESA Rule Deployment Name or Delete a Deployment	127
Edit the ESA Rule Deployment Name	128
Delete an ESA Rule Deployment	128
Show Updates to an ESA Rule Deployment	129
Fast Deployment and Full Deployment of ESA Rules	130
View ESA Stats and Alerts	133
View Stats for an ESA Service	133
View ESA Stats	133
Enable or Disable Rules	134
Refresh the Statistics	134
View a Summary of Alerts	135

Add an Advanced EPL Rule	138
Prerequisites	138
Add an Advanced EPL Rule	139
Validate an Advanced EPL Rule	141
Event Processing Language (EPL)	146
ESA Annotations	148
@RSAContext Annotation (11.5 and later)	148
Prerequisites	148
Single-Column Context Hub Lists	148
Multi-Column Context Hub Lists	149
Automatic Context Hub List Updates	151
Single Column Context Hub List Update Example	151
Multi-Column Context Hub List Update Example	151
@RSAAlert Annotation	151
@RSAPersist Annotation	153
@UsesEnrichment (10.6.1.1 and later)	153
@Name	153
@Audit Annotation	154
Example Advanced EPL Rules	155
Example #1:	155
EPL #1:	155
EPL #2:	156
Example #2:	156
EPL #3:	156
EPL #4: Using NamedWindows and match recognize	157
EPL #5: Using Every @RSAAlert(identifiers={"user_src"})	157
Example #3:	158
EPL #6: @RSAAlert(identifiers={"ip_src"})	158
EPL #7: @RSAAlert(identifiers={"ip_src"})	159
Example #4:	159
EPL #8: using time_batch	159
Example #5:	160
EPL #9: using timer:interval	160
Example #6:	160
EPL #10: using timer and Lockout	160
Example #7:	161
EPL #11: @RSAAlert(oneInSeconds=0)	161
Configure an In-Memory Table Using an EPL Query	163
Workflow	163
Prerequisites	164

Procedure	164
Example	165
Step 1: Create the Enrichment	165
Step 2: Create Your Rule	167
Rule Statement	167
Rule Logic with Enrichment Added	168
ESA Alert References	169
Rules Tab	170
What do you want to do?	170
Related Topics	170
Quick Look	171
Rules Tab Options Panel	172
What do you want to do?	172
Related Topics	172
Quick Look	172
Rules Section	172
Deployments Section	172
Rule Library Panel	174
What do you want to do?	174
Related Topics	174
Quick Look	174
Rule Library Toolbar	175
Rule Library List	175
Rule Builder Tab	177
What do you want to do?	177
Related Topics	177
Quick Look	177
Conditions Section	180
Notifications Section	182
Enrichments Section	182
Debug Option	183
Test Rule Section	183
Syntax	185
Build a Statement Dialog	186
What do you want to do?	186
Related Topics	186
Quick Look	186
Advanced EPL Rule Tab	191
What do you want to do?	191
Related Topics	191

Quick Look	191
Notifications Section	194
Enrichments Section	195
Test Rule Section	195
Syntax	197
Rule Syntax Dialog	198
Quick Look	198
Deployment Panel	200
What do you want to do?	200
Related Topics	200
Quick Look	201
ESA Services Section	201
Data Sources Section	202
Position Tracking Information	203
Editing Position Tracking Information	203
Use Case Scenario	205
Importing Position Tracking Information	206
Deployment Options	207
ESA Rules Section	208
Data Source Filter (Optional) Section	209
Create Data Source Filter Dialog - Simple	210
Create Data Source Filter - Advanced	211
Deploy ESA Services Dialog	213
What do you want to do?	213
Related Topics	213
Quick Look	213
Deploy ESA Rules Dialog	215
What do you want to do?	215
Related Topics	215
Quick Look	215
Updates to the Deployment Dialog	217
What do you want to do?	217
Related Topics	217
Quick Look	217
Services Tab	219
What do you want to do?	219
Related Topics	219
Quick Look	219
ESA Services Panel	220
General Stats Panel	220

Deployed Rule Stats Panel	221
Settings Tab	223
What do you want to do?	223
Related Topics	223
Quick Look	223
Meta Key References	224
Viewing the List of Meta Entities	224
Enabling Meta Entity in the ESA Correlation Server	224
Meta Entity Usage Example	225
Building Rules with Custom Meta Entities	226
Prerequisites	226
Enrichment Sources	226

Getting Started with ESA

This topic covers quick start topics for NetWitness Event Stream Analysis (ESA) to help you get started in using ESA. The following topics are designed to assist you in working with ESA Correlation Rules.

- [Best Practices](#) helps you to understand how to best set up, deploy, and create rules.
- [Troubleshoot ESA](#) helps you to troubleshoot different aspects of ESA, including rule writing and deployment.
- [View Memory Metrics for Rules](#) helps you to work with memory metrics to understand memory usage for ESA services.

In NetWitness version 11.5 and later, There are only two services that can run on an ESA host:

- **ESA Correlation (ESA Correlation rules):** Creates alerts from ESA rules.
- **Contexthub Server (Context Hub):** Runs only on an ESA primary host. Contexthub Server provides enrichment lookup capability in the Respond and Investigate views. For information, see the *Context Hub Configuration Guide*.

Note: The Event Stream Analytics Server (ESA Analytics) service is not supported in NetWitness Platform version 11.5 and later.

The first service is the ESA Correlation service that creates alerts from ESA rules, also known as ESA Correlation Rules, which you create manually or download from Live.

In NetWitness 11.3 and later, the ESA Correlation service replaces the Event Stream Analysis service and is also known as ESA Correlation Server. The ESA Correlation service provides the same services as the Event Stream Analysis service with the added benefit of enabling you to specify different data sources for your ESA correlation rules. Like the Event Stream Analysis service, the ESA Correlation service installs on the ESA Primary and ESA Secondary host types.

The second service is the Contexthub Server service, which provides enrichment lookup capabilities in the Respond and Investigate views. It runs only on an ESA Primary host. For information, see the *Context Hub Configuration Guide*. Go to the [NetWitness Documents - All Versions](#) to find all RSA NetWitness Platform 11.x documents.

IMPORTANT: The NetWitness server (Admin server), ESA primary host, and ESA secondary host must all be on the same NetWitness Platform version.

How ESA Generates Alerts

The ESA Correlation service runs rules that specify criteria for problem behavior or threatening events in your network. When ESA detects a threat that matches rule criteria, it generates an alert.

To generate alerts, ESA performs the following functions:

1. Gathers data
2. Runs ESA rules against the data
3. Captures events that meet rule criteria
4. Generates alerts for those captured events

Data Source Configuration Changes

In NetWitness version 11.3 and later, the ESA Correlation service enables you to specify different data sources for different sets of rules. Instead of adding data sources, such as Concentrators, to the entire ESA Correlation service, you can specify different data sources for each ESA rule deployment. An ESA rule deployment includes an ESA Correlation service with its associated data sources and a set of ESA rules. For example, you may want to use Concentrators with HTTP packet data in one deployment and Concentrators with HTTP log data in another deployment. For more detailed information, see [Deploy Rules to Run on ESA](#).

An Endpoint Risk Scoring Rules Bundle is available in NetWitness Platform

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

The ESA Correlation service can process endpoint risk scoring rules, which generate alerts that are used in risk scoring calculations to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see "Deploy Endpoint Risk Scoring Rules on ESA" in the *ESA Configuration Guide*. For complete information on configuring NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*.

Go to the [NetWitness Documents - All Versions](#) to find all RSA NetWitness Platform 11.x documents.

Best Practices

Best practices provide guidelines to help you write and manage rules, deploy rules, and maintain system health for your ESA services.

Understand Event Stream Analysis Rule Types

The ESA Correlation service provides advanced stream analytics such as correlation and complex event processing at high throughputs and low latency. It is capable of processing large volumes of disparate event data from Concentrators. However, when working with Event Stream Analysis, you should be aware of the factors that affect resource usage in order to create effective rules.

Each event that is received by ESA is evaluated to determine if it may trigger a rule. There are three types of rules that can be deployed in order to determine what the ESA engine should do with the incoming event. Each of these rule types have different impacts on system resource utilization. All three rule types may be created via the Rule Builder, Advanced Event Processing Language (EPL) rules, or downloaded via RSA Live. The table below lists the rule type and the impact this rule may have on system resources.

Rule Type	Description
Simple Filter Rule	<p>This rule has no correlation to other events. At ingestion time, this rule is evaluated against a set of conditions, and if those conditions are met an alert is generated. If no conditions match, the event is quickly released by the engine to free up memory usage. These rules do not take up memory since the events are not retained beyond the initial evaluation. The memory resource usage does not increase as more simple filter rules are deployed. However, if the filter condition is too generic, it is possible that this rule can generate too many alerts, which will strain the system resources for the storage and retrieval of these alerts.</p> <p>For example, you might write a rule to generate an alert when HTTP network activity arrives over a non-standard HTTP port.</p>
Event Window Rule	<p>This rule evaluates a set of events over a time period for specific conditions. At ingestion time, the rule is evaluated against a set of conditions. If those conditions are met, the event is retained in memory for a specific amount of time. After the specified time passes, the events are removed from the time window if the number of events collected does not meet the threshold to trigger an alert.</p> <p>The memory consumption of such rules is highly dependent on the incoming event rate (traffic), the amount of data per event, and the time length specified in the event window. Each matching event is retained in memory until the time window has passed, so the longer the time window, the greater the potential volume. For example, you might write a rule that generates an alert if a user has five failed login attempts within a ten minute time frame.</p>

Rule Type	Description
Followed By Rule	<p>This rule evaluates a chain of incoming events to determine if the sequence of events matches a particular condition. At ingestion time, the rule is evaluated against a set of conditions. If the conditions are met, one of two actions occurs:</p> <ul style="list-style-type: none"> • If this is the first event of the sequence, a new event thread is started, and the event is retained as the head of the sequence. • If the event belongs to an existing event thread, it is added to that sequence. <p>In both cases, the event is retained in memory. The amount of resource usage is particularly sensitive to the customer environment for this type of rule. If the filter condition generates many event threads, resources are consumed for each new thread (in addition to the event). Additionally, if the end of the event thread is never met (that is, an alert is never generated), then the entire event is saved in memory indefinitely. For example, you might write a rule to generate an alert when a user fails to log in to a server, then performs a successful login, and then creates a new account.</p>


Note: ESA sends alerts to NetWitness Respond for processing and the alerts are eventually stored in a database. If your rule creates too many alerts, it can slow down another part of the system.

When writing and deploying rules, you should be aware that rule memory usage and alert generation consume system resources. The sections below are designed to help you keep your usage at a healthy level and monitor for problems if systems are becoming overloaded.

Best Practices for Writing Rules

These are general guidelines for writing rules.

- **Create alerts for actionable events.** The purpose of an alert should be to notify you of an event that requires immediate and specific action. For events that do not require action, or only require you to have awareness of the event, you can create a report.
- **Validate your ESA rules within the Rule Builder or Advanced EPL Rule Builder before you deploy them.** To prevent errors in your rules and confirm that they generate the expected alerts, you can test the rule logic with JSON data within the rule builders. This capability is available in NetWitness Platform version 11.5 and later. For more information, see [Validate an ESA Rule](#) and [Validate an Advanced EPL Rule](#).
- **Configure new rules as trial rules so you can observe how they react in your environment.** If you deploy new rules as trial rules, they will be disabled if the configured memory threshold is exceeded.
- **Configure Alert notifications only after your rule testing and tuning is complete.** This can help ensure you do not get flooded with notifications if a rule behaves differently than you expect.

- **Rules need to be specific so that you limit resource usage.** Use the following guidelines to limit usage:
 - Make the filters on the rule exclude all but the necessary events for the rule to fire accurately.
 - Make the size of your windows (window time for correlation) as small as possible.
 - Limit the events that you include in the window: For example, if you only want to see IDS events, ensure that you only include those events in your time window.
- **Add Memory Thresholds to ESA rules that use memory.** For example, if a rule contains windows or pattern matching, configure a memory threshold for that rule. If the rule goes over the allotted memory threshold, it gets disabled individually and an error is displayed for that rule on the  **(Configure) > ESA Rules > Services** tab. This capability is available in NetWitness Platform version 11.5 and later.
- **Rules need to be tuned to an alert level that is manageable.** If you are flooded with alerts, then the purpose and of an alert is lost. For example, maybe you want to know about encrypted traffic to other countries. But, you could limit the list to countries that are known risks. This limits the volume of alerts to a level you can manage.

For more best practice information for writing ESA rules, see [ESA Rule Writing Best Practices](#).

Best Practices for Working with RSA Live Rules

These are guidelines for RSA Live Rules.

- **Deploy RSA Live rules in small batches.** Not every rule is suited to every environment. The best way to ensure your RSA Live rules are successful is to deploy them in small batches so you can test them in your environment. If you deploy small batches, it's much easier to tell if a particular rule has an issue.
- **Read the rule descriptions provided with RSA Live rules.** ESA rules are not “one size fits all.” Not all rules will work in your environment. The rule descriptions tell you which parameters you will need to modify to successfully deploy a rule in your environment.
- **Set your parameters.** RSA Live rules have parameters that need to be modified. If you do not modify your parameters, the rule may not work or it may exhaust your memory.
- **Deploy new rules as trial rules so you can observe how they react in your environment.** If you deploy new rules as trial rules, they will be disabled if the configured memory threshold is exceeded. For more details, see [Working with Trial Rules](#).

Best Practices for Deploying Rules

These are general guidelines for deploying rules.

- **Deploy rules in small batches so you can observe how they react in your environment.** Not all environments are the same, and a rule will need to be tuned for memory usage, alert volume, and

effective detection of events.

- **Test rules before you configure alert notifications.** Configure Alert notifications only after your rule testing and tuning is complete. This can help ensure you do not get flooded with alerts if a rule behaves differently than you expect.
- **Monitor system health as a part of your deployment process.** When you deploy rules, monitor your system's health as a part of your deployment process. You can view total memory usage for your ESA in the Health and Wellness tab. For more information, see "View Detailed Statistics in Health and Wellness" in [Troubleshoot ESA](#).

Best Practices for System Health




These are general guidelines for system health.








- **Set up new rules as trial rules.** A common issue is that new rules may cause memory issues. To prevent this, you can set up new rules as trial rules. If the configured memory threshold is met, all trial rules are disabled to prevent the system from running out of memory. For more information about trial rules, see [Working with Trial Rules](#).
- **Set up thresholds in Health & Wellness to alert you if memory usage is too high.** There are metrics in NetWitness Health & Wellness that track memory usage. You can set up alerts and notifications to send you an email if those thresholds are crossed. For more information about the memory statistics you can view, see [View Memory Metrics for Rules](#).
- **Monitor memory metrics for each rule in Health & Wellness.** For each rule, you can view the estimated memory usage in Health & Wellness. You can use this information to ensure that rules do not use too much memory. For more information about the memory statistics you can view, see [View Memory Metrics for Rules](#).

Troubleshoot ESA


This section describes common issues that may occur while using ESA, and it suggests common solutions to these problems.

Troubleshoot ESA Correlation Services



Problem	Possible Causes	Solutions
<p>On the NetWitness Platform Dashboard, the ESA service appears in red to indicate it is offline.</p> <p>In the  (Configure) > ESA Rules view, the following message appears: "The Service is either offline or not reachable."</p>	Several	<p>When an ESA Correlation service is offline, there are many possible causes. However, a common issue is that you have created a rule that uses excessive memory and causes the ESA service to fail. To troubleshoot this problem, see Steps to Troubleshoot Memory Issues with an ESA Service Offline.</p> <p>Other common causes might be that your firewall is blocking the connection between the ESA and NetWitness Platform, or the ESA Correlation service machine may be down.</p> <p>To bring up ESA Services:</p> <p>Go to  (Admin) > Services, select your ESA service, and then select  > Start.</p> <p>If your ESA service is stopping and restarting in a loop, you may need to call Customer Support to get the services to start.</p>


Problem	Possible Causes	Solutions
<p>On the NetWitness Platform Dashboard, the ESA service appears in red to indicate it is offline.</p> <p>In the  (Configure) > ESA Rules view, the following message appears: "The Service is either offline or not reachable."</p>	<p>Configuration issues</p>	<p>If your system has been recently upgraded, you may have made a configuration error.</p> <p>Go to  (Admin) > Services, select your ESA service and then select   > Edit. In the Edit Service dialog, click Test Connection. If the connection fails, you likely have a configuration error. Attempt to fix your configuration error and try again.</p>
<p>Suddenly, an ESA Correlation service is completely unresponsive and appears to have crashed. The ESA Correlation Log file shows a Too Many Open Files error.</p>	<p>Connectivity issues</p>	<p>The most likely cause for this error is a connectivity issue between the ESA Correlation service and a data source used in an ESA rule deployment, such as a Concentrator or Decoder.</p> <ol style="list-style-type: none"> Restart the ESA Correlation service. Go to  (Admin) > Services, select your ESA service and then select   > Restart. Check the connectivity from ESA to the data source. Look for connection errors in the ESA Correlation logs. You can use SSH to get in the system and go to: <pre>/var/log/netwitness/correlation-server/correlation-server.log.</pre> <p>For example:</p> <pre>Error: com.rsa.netwitness.streams. RecordStreamException: admin@<ip address>:56005:: com.rsa.netwitness.streams.RecordStreamException: connect::com.rsa.asoc.transport.nw.session. NextgenException: Failed to connect to <ip address>:56005</pre> <p>If there are network connectivity issues, fix the issues and then restart the ESA Correlation service again to see if it fixes the problem.</p> <p>If you have a data source with intermittent connectivity, you should remove it from the ESA rule deployment.</p>



Troubleshoot RSA Live Rules for ESA



Problem	Possible Causes	Solutions
<p>I imported a group of rules from RSA Live, and now my ESA service is crashing. Why?</p>	<p>You may not have configured the parameters for the RSA Live rule to tune it for your environment.</p>	<p>Each rule in RSA Live has a description that includes the parameters you must configure and prerequisites for your environment. Review this description to see if the rule is appropriate for your environment.</p> <p>To ensure that you deploy rules safely in your environment, configure new rules as trial rules to test them in your environment. Trial rules add a safeguard for testing new rules. For details on this, see Deploy Rules as Trial Rules.</p>
<p>I imported a group of rules from RSA Live, and while the rules deployed without errors, they were later disabled.</p>	<p>Not all RSA Live rules are meant for every environment. You may not have the correct meta in your ESA for the rule to run.</p>	<p>You can verify that a rule was disabled by going to  (Configure) > ESA Rules > Services > Deployed Rule Stats. If the rule is disabled, the green icon does not display next to the rule.</p> <p>If a rule deployed correctly but was disabled, check the logs for exceptions related to the rule. Specifically, check to see if the rules were disabled due to missing meta. To do this, go to the ESA Correlation logs. You can use SSH to get in the system and go to:</p> <pre data-bbox="711 1018 1221 1075">/var/log/netwitness/correlation-server/correlation-server.log.</pre> <p>Then, search for a message similar to the following:</p> <pre data-bbox="711 1138 1383 1194">"Property named '<meta_name>' is not valid in any stream"</pre> <p>For example, you might see:</p> <pre data-bbox="711 1255 1399 1360">Failed to validate filter expression '(medium=1 and streams=2 or medium=3...(238 chars)': Property named 'tcp_flags_seen' is not valid in any stream</pre> <p>If a similar message displays, you may need to add a custom meta key to the Log Decoder or Concentrator. To do this, follow these instructions: "Create Custom Meta Keys Using Custom Feed" in the <i>Decoder and Log Decoder Configuration Guide</i>.</p>

Troubleshoot ESA Rules

Problem	Possible Causes	Solutions
<p>I have an ESA rule that is not getting deployed and is not creating alerts.</p>	<p>A meta key that the rule uses is a string array type, but it shows as a string type on ESA.</p>	<p>Check to see if any string array meta keys that the rule uses are configured as string array types on ESA. Go to  (Configure) > ESA Rules > Settings tab (Meta Key References).</p> <ul style="list-style-type: none"> • If it shows <code>string[]</code>, it is configured as a string array type on ESA. This is fine. • If it shows <code>string</code> without the brackets, it is configured as a string type and you need to fix it on ESA. <p>In the ESA Correlation service Explore view, go to <code>correlation/stream</code>. Add string array meta keys to the multi-valued list to allow them to be used as an array in ESA rules. Go back to the Meta Key References and click the refresh icon (). Verify that the meta keys with a string array type show a value of <code>.string[]</code>. For additional details, see "Configure Meta Keys as Arrays in ESA Correlation Rules" in the <i>ESA Configuration Guide</i>.</p>

Problem	Possible Causes	Solutions
<p>I created a rule, and I checked the syntax. The rule looked fine. When I went to deploy the rule, I got an error. Why?</p>	<p>You may not have the correct meta to deploy the rule.</p>	<p>Check the Meta Key References. You may not have the correct meta to deploy the rule. Check the ESA Correlation service log files to see which meta keys are missing: <code>/var/log/netwitness/correlation-server/correlation-server.log</code></p> <p>In NetWitness version 11.3.0.2 and later, you can check the ESA rule status in the ESA rule deployment (Go to  (Configure) > ESA Rules > Rules tab. In the options panel on the left, select a deployment and look in the ESA Rules section). If a disabled rule has an error message, it shows  in the Status field. Hover over the rule to view the error message tooltip.</p>  <p>In the above example, the error message shows: "Failed to resolve event type: Event type or class named "Host_Whitelist" was not found." In this case, a Context Hub list called "Host_Whitelist" that is used by the rule is not available. For more information on context hub lists, see the <i>Context Hub Configuration Guide</i>.</p> <p>For more information, see the ESA Rules Section section of the <i>Deployment Panel</i> reference.</p>
<p>I set up notifications for a rule, but we are not receiving them. The correlation-server.log file does not show any errors. Why?</p>	<p>Correlation-server successfully sent the notification messages to integration-server, but when integration-server tried to send the notifications to their destination, it failed.</p>	<p>When troubleshooting notifications, check both the ESA Correlation service log files (<code>/var/log/netwitness/correlation-server/correlation-server.log</code>) AND the Integration-Server log files on the NetWitness Server (<code>/var/log/netwitness/integration-server/integration-server.log</code>). For an example, see Integration-Server SMTP Notification Error Example.</p> <div style="border: 1px solid green; padding: 5px;"> <p>Note: For any notification-related troubleshooting, check the integration-server log file in addition to the log file of the service creating the notification.</p> </div>

Problem	Possible Causes	Solutions
<p>I created a rule with an enrichment, added an SMTP notification, and deployed my rule. We are not receiving SMTP notifications. Why?</p>	<p>You do not have a template that met the criteria to parse the events.</p>	<p>Check the ESA Correlation service log files to see if the SMTP notification failed: <code>/var/log/netwitness/correlation-server/correlation-server.log</code>. For more details on the notification error, check the Integration-Server log file on the NetWitness Server (also known as Node 0, Admin server, or NWServer): <code>/var/log/netwitness/integration-server/integration-server.log</code>.</p> <p>If you use an ESA rule that has an enrichment, such as a Context Hub list, you must create a custom template. You can duplicate a default template and adjust it for your enrichment. See SMTP Notification Error Example below for a notification error example.</p> <p>For information on creating a custom ESA template, see "Define a Template for ESA Alert Notifications" in the <i>System Configuration Guide</i>.</p> <p>Go to the NetWitness Documents - All Versions to find all RSA NetWitness Platform 11.x documents.</p>
<p>I created a custom rule (via the Rule Builder or Advanced EPL), and my rule is not firing. Why?</p>	<p>You may have connectivity issues.</p>	<p>Check the Offered Rate statistic on the  (Configure) > ESA Rules > Services tab. Select the ESA service and then look at the statistics on the tab for the Deployment.</p> <p>If the Offered Rate is zero, then the ESA service is not receiving data from Concentrators. Check the ESA Correlation log files for connectivity issues: <code>/var/log/netwitness/correlation-server/correlation-server.log</code>.</p> <p>If the offered rate is not zero, the meta key name and type used in the rule likely doesn't match the meta key present in events. Check to see if the meta key name and type used in the rule is valid by searching for the meta key name in  (Configure) > ESA Rules > Settings tab (Meta Key References).</p>

Problem	Possible Causes	Solutions
I created a custom rule (via the Rule Builder or Advanced EPL), and my rule is not firing. Why?	There may be a problem with the rule.	<p>If a specific rule is not firing, go to  (Configure) > ESA Rules > Services to see if the rule was disabled. In the Deployed Rule Stats section, a rule that is disabled displays a clear enabled button (instead of the green enabled button).</p> <p>You can also check Events Matched field. Go to  (Configure) > ESA Rules > Services. From there, you can see the number of events that were matched in the Events Matched column.</p> <p>If no events matched, check the logic of your rule for errors. For example, check the syntax for uppercase and lowercase errors, and check the time window. If the rule still doesn't fire, consider simplifying the logic of the rule to see if it fires when there is less complexity.</p>
After a recent upgrade, I am not seeing alerts and I am seeing disabled rules.	There may be a problem with the ESA rule deployment.	<p>Deploy the ESA rule deployments again. ESA Rule Deployment Steps provides more information on deploying rules using the ESA Correlation service.</p> <p>If this does not resolve the issue, check the ESA Correlation log files for more information: <code>/var/log/netwitness/correlation-server/correlation-server.log</code>.</p>
After an update or upgrade to 11.3.0.2 or later, if I try to make an adjustment to some rules, I get an error when trying to save them.	The Ignore Case option may be selected for a meta key that does not contain alphabetic values, such as IP address.	<p>In NetWitness Platform 11.3.0.2 and later, the Ignore Case option has been removed from the ESA Rule Builder - Build a Statement dialog for meta keys that do not contain text values. Adding Ignore Case on meta keys which do not contain alphabetic values causes additional processing to occur for no added benefit.</p> <p>In the ESA Rule Builder - Build a Statement dialog, check to see if you have any meta keys that do not contain alphabetic characters, for example, ip_src and ip_dst. If you do, clear the Ignore Case checkbox for those meta keys and try to save the rule again.</p>

Problem	Possible Causes	Solutions
<p>After an upgrade to 11.3.0.2 or later, I see a warning message in the ESA Correlation service log file showing a difference between the multi-valued and default-multi-valued parameter meta key values. Why?</p>	<p>You do not have the required meta keys on ESA Correlation that the Endpoint, UEBA, and Live content rules need to work.</p>	<p>If you want to use the latest Endpoint, UEBA, and Live content rules, add the necessary meta keys to the multi-valued and single-valued parameter fields. For detailed information and instructions, see "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the <i>ESA Configuration Guide</i>.</p> <p>For example warning messages, see Example ESA Correlation Server Warning Message for Missing Meta Keys.</p>
<p>Meta keys marked as sensitive for data privacy are still included in notifications and alerts for some rules.</p>	<p>In ESA rules that do not select every piece of meta from the session (that is, using select *), you may see that data privacy (if enabled) and the Pivot to Investigate > Navigate link accessed from a context tooltip in the Respond Incident Details view does not work.</p>	<p>The following steps apply to all released versions of NetWitness 11.3 and later. In 11.4 and later, you do not need to follow these steps for data privacy. However, you need to follow these steps if you want to enable the Pivot to Investigate > Navigate link accessed from a context tooltip in the Respond Incident Details view.</p> <ol style="list-style-type: none"> 1. Add the ESA generated <code>event_source_id</code> meta key to the <code>index-concentrator-custom.xml</code> file. 2. Add the <code>event_source_id</code> meta key to the SELECT statement within any ESA rule that does not select every piece of metadata from the session. 3. After the Concentrator changes take effect, redeploy the ESA rule deployment that contains the ESA rule. <p>To do this, see Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_source_id.</p> <p>For NetWitness 11.4 and later, to resolve the data privacy issue, see How to Remove Sensitive Meta Keys Globally from All Alerts.</p>

Problem	Possible Causes	Solutions
<p>The Pivot to Investigate > Navigate link does not work in a context tooltip accessed from Respond.</p>	<p>In ESA rules that do not select every piece of meta from the session (that is, using select *), you may see that data privacy (if enabled) and the Pivot to Investigate > Navigate link accessed from a context tooltip in the Respond Incident Details view does not work.</p>	<p>For NetWitness 11.3 and later (including 11.4), see Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event source id.</p>
<p>I added a data source filter to the data sources in my ESA rule deployment. It was working fine and then all of a sudden, I stopped receiving alerts.</p>	<p>If there are any adjustments to an application rule on the Decoders that are mapped to the data sources used in your data source filter, the filter does not work for that rule since the application rule used in the filter no longer exists.</p>	<p>The optional data source filter is available in NetWitness 11.5 and later.</p> <p>If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.</p>

Problem	Possible Causes	Solutions
<p>I added a data source filter to the data sources in my ESA rule deployment and I see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file.</p>	<p>You are filtering out a large portion of the traffic and less sessions match the aggregation criteria.</p>	<p>In the Explore view node list for an ESA Correlation service, select correlation > stream. Decrease the <code>max-sessions</code> parameter from 10,000 to a lower session count and restart the ESA Correlation service. For step-by-step instructions, see Adjust Maximum Sessions for the Data Source Filter.</p>
<p>I created a rule using a Context Hub list and it was working properly, but suddenly, the rule stopped firing and ESA is not processing any rules.</p>	<p>If a Context Hub list that is used by ESA Correlation is renamed or deleted, ESA will not be able to access the list and may halt processing for all rules.</p>	<p>Do not rename or delete a Context Hub list that is used in a deployed ESA rule. ESA Correlation will also not be able to access a Context Hub list if you delete it and then add it again with the same name while the rule is deployed.</p> <p>If you rename a Context Hub list or recreate the Context Hub list with the same name, update the ESA rules that use that Context Hub list, and then redeploy the ESA rule deployments that contain those rules.</p>

SMTP Notification Error Example


The following SMTP notification error example is an excerpt from a `correlation-server.log` file, which shows an error message for sending notifications with unsupported templates. In this example, there is a rule that is configured with the GeoIP enrichment, which has a hash table as one of its fields (the GeoIPLookup meta). Because the default SMTP template is only designed to deal with metas that are either singular values or arrays that contain only singular values, such as `"ip.src": "1.1.1.1"` and `"action": ["fw:inbound-network-traffic"]`, sending the email notification fails due to the array containing a hash table.

FTL stack trace ("~" means nesting-related):

```
- Failed at: ${value!""} [in template "smtp.ftl" in macro "value_of" at line 1, column 152]
- Reached through: @value_of metadata[key] [in template "smtp.ftl" at line 85, column 141]
----
...
For "${...}" content: Expected a string or something automatically convertible to string (number, date or boolean), or "template output" , but this has evaluated to an extended_hash (LinkedHashMap wrapped into f.t.DefaultMapAdapter):
==> value!"" [in template "smtp.ftl" at line 1, column 154]
```

Integration-Server SMTP Notification Error Example

The following SMTP notification error example is an excerpt from an `integration-server.log` file, which shows a failure when the Integration-server attempts to send an email notification to the email notification server. In this case, you should check the email notification server configuration in the

Global Notifications settings ( (Admin) > System > Global Notifications > Servers tab).

```
2019-10-09 18:53:42,015 [-SMTP-5c45c867e4b03b89a49b78ba] WARN
Notification|SMTP dispatch failed (Reason: Sending the email to the following
server failed : email.server.com:25)
2019-10-09 18:53:42,100 [-SMTP-5c45c867e4b03b89a49b78ba] WARN
SystemOperation|Failed to forward ResolvedNotification
{server=5c45c867e4b03b89a49b78ba, destination=5c45c854e4b03b89a49b78b9,
content-length=30681}
java.lang.IllegalArgumentException: org.apache.commons.mail.EmailException:
Sending the email to the following server failed : email.server.com:25
```

Example ESA Correlation Server Warning Message for Missing Meta Keys

If you see a warning message in the ESA Correlation server error logs for missing multi-valued meta keys, there is a difference between the `default-multi-valued parameter` and `multi-valued parameter` meta key values, and the new Endpoint, UEBA, and Live content rules will not work. The same is true for missing single-valued meta keys. Completing the "Update the Multi-Valued and Single-Valued Parameter Meta Keys for the latest Endpoint, UEBA, and RSA Live Content Rules" procedure in the *ESA Configuration Guide* should fix the issue.

Multi-Valued Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[alert, alert_id, browserprint, cert_thumbprint, checksum, checksum_all, checksum_dst, checksum_src, client_all, content, context, context_all, context_dst, context_src, dir_path, dir_path_dst, dir_path_src, directory, directory_all, directory_dst, directory_src, email_dst, email_src, feed_category, feed_desc, feed_name, file_cat, file_cat_dst, file_cat_src, filename_dst, filename_src, filter, function, host_all, host_dst, host_orig, host_src, host_state, ip_orig, ipv6_orig, OS, param, param_dst, param_src, registry_key, registry_value, risk, risk_info, risk_suspicious, risk_warning, threat_category, threat_desc, threat_source, user_agent] are still MISSING from multi-valued
```

Single Value Warning Message Example

```
2019-08-23 08:55:07,602 [ deployment-0] WARN Stream|[accesses, context_target, file_attributes, logon_type_desc, packets] are still MISSING from single-valued
```

Update any ESA Rule that Selects Only Certain Meta Keys from the Session to Include event_source_id

In ESA rules that do not select every piece of meta from the session (that is, using `select *`), you may see that data privacy (if enabled) and the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view does not work.

The following steps apply to all released versions of NetWitness 11.3 and later. In 11.4 and later, you do not need to follow these steps for data privacy, instead, see [How to Remove Sensitive Meta Keys Globally from All Alerts](#). However, you need to follow these steps if you want to enable the **Pivot to Investigate > Navigate** link accessed from a context tooltip in the Respond Incident Details view.

Note: Do not use any Esper keyword as custom meta keys since this causes an error while creating an ESA Rule. For Esper keywords, see [Reserved keywords](#).

1. Update the `index-concentrator-custom.xml` file to include the ESA generated `event_source_id` meta key. If you do not add the meta key, ESA cannot recognize it and the rule will fail to deploy.

The following figure shows the file configured for the custom meta key "Event Source ID" with index settings of "IndexNone" with a format of "Text".

```
<key description="Event Source ID" name="event_source_id" format="Text" level="IndexNone"/>
```

(decoder/logdecoder) will be transformed and the resulting value persisted in another key, informational when specified on other services
destination = specifies the key name of the transformed meta value to create

Decoder examples - Normally you do not need to edit index files on the Decoder, unless you want to add aliases or have data privacy requirements. Parsers and feeds declare their meta keys internally and those keys are automatically added to the language. Also, you should *never* set the index level to IndexKeys or IndexValues on a Decoder if you have a Concentrator/Archiver aggregating from it. The index partition size is too small to support any indexing beyond the default "time" meta.

Data privacy
 <key description="existing meta key" format="Text" level="IndexNone" name="existing" protected="true">
 <transform destination="existing.hash"/>
 </key>

Concentrator/Archiver examples - Any new meta keys that should be indexed must be added to this file.

Adding new meta key for custom parser at the index key level
 <key description="my new parser meta key" format="Text" level="IndexKeys" name="mynewparserkey"/>

Data privacy
 <key description="existing meta key" format="Text" level="IndexValues" name="existing" protected="true">
 <transform destination="existing.hash"/>
 </key>
 <key description="existing meta key hash" format="Text" level="IndexValues" name="existing.hash" token="true"/>

Broker derives its language from all the devices it aggregates from. There is simply no need to edit a broker's custom language file.
 -->

<!-- *** Please insert your custom keys or modifications below this line *** -->
 <key description="Event Source ID" name="event_source_id" format="Text" level="IndexNone"/>
 </language>

To save and deploy the new setting on the NetWitness host, select the **Apply** button. To force the change, restart the Concentrator service or you can wait until the next polling interval for the change to be recognized.

The XML file can also be deployed to other NetWitness hosts by clicking the **Push** button and selecting the destination NetWitness host. Only deploy the XML file to a NetWitness host that runs that service (that is, other Concentrators).

- Update any ESA rule that selects only certain meta from the session to include the ESA generated **event_source_id** meta key. Add the **event_source_id** meta key to the **SELECT** statement. See the following example rule.

Example Rule with event source id

```
@RSAAlert
SELECT user_dst, reference_id, hostname, event_source_id FROM
Event(
```

```

device_class='Windows Hosts',
reference_id IN ('4624' , '4625'),
user_dst IS NOT NULL,
user_dst NOT LIKE '%$ ',
user_dst NOT IN ('ANONYMOUS LOGON','SYSTEM')
)
.win:time_batch(10 Minutes)
GROUP BY user_dst
HAVING COUNT(distinct hostname) >= 15;


```


3. After the Concentrator changes takes effect, redeploy the ESA rule deployment that contains the ESA rule.

For additional information, see [How ESA Handles Sensitive Data](#). For information on the strategy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

Steps to Troubleshoot Memory Issues with an ESA Service Offline

Step 1: Verify that your Host Is Running

The first step to troubleshooting is to ensure that your host is running. To do this, go to  (Admin) > **Hosts**. If the host is down, the system parameters will not display (updating host information can sometimes be delayed), the **Services** display in red, and you may see an error message.





Name	Host	IP	Services	Current Version	Update Version	Status
ELH	10.10.10.01	10.10.10.01	4	11.5.0.0		Up-to-Date
ESA	10.10.10.02	10.10.10.02	2	11.5.0.0		Up-to-Date
PH	10.10.10.03	10.10.10.03	2	11.5.0.0		Up-to-Date
SA	10.10.10.04	10.10.10.04	12	11.5.0.0		Up-to-Date
UEBA	10.10.10.06	10.10.10.06	1	11.5.0.0		Up-to-Date

If your host is down, contact your NetWitness Administrator to restart it. Otherwise, go to Step 2.


Step 2: View Detailed Statistics in Health & Wellness

If your ESA service is down, you can go to Health & Wellness and view the **last known metrics** to see where potential issues are occurring. The most common problem is that your ESA service is exceeding memory thresholds, which causes it to stop or fail. In NetWitness 11.5 and later, see also [View Health Statistics and Trends for ESA Correlation in New Health & Wellness](#).

- Go to  (Admin) > **Health & Wellness** > **Alarms** to see if the ESA triggered any alarms. Look for the following alarms for ESA Correlation:
 - Correlation Server in Critical State
 - Correlation Server in Unhealthy State
 - Correlation Server Stopped




Time	State	Severity	Rule Name	Service	Hostname	IP Address	Stat
2020-05-21 11:34:05 PM	Active	High	Power Supply Failure	Host	CoreWithDAC	10...	Power Supply Status 0x62
2020-05-21 01:50:13 PM	Active	High	Power Supply Failure	Host	HybridEPH2	10...	Power Supply Status 0x62
2020-05-21 07:49:56 AM	Active	High	Meta(s) Overflow in Index Slices	Concentrator	NetworkHybrid	10...	Index/Meta(s) Overflow
2020-05-21 05:45:58 AM	Active	High	Custom Feeds Failure	NetWitness UI	SA	10...	Feeds/Custom Feeds Deployment Status
2020-05-21 05:35:01 AM	Active	High	Concentrator Not Consuming From Service	Concentrator	EPH	10...	
2020-05-20 12:56:23 PM	Active	High	Power Supply Failure	Host	Core2	10...	
2020-05-20 12:55:25 PM	Active	High	Power Supply Failure	Host	Core	10...	Power Supply Status 0x63
2020-05-19 06:58:12 PM	Active	High	Broker Session Rate Zero	Broker	Malware	10...	Broker/Session Rate (current)
2020-05-19 06:54:14 PM	Active	High	Power Supply Failure	Host	S45Malware	10...	Power Supply Status 0x62
2020-05-19 03:08:59 PM	Active	High	Power Supply Failure	Host	UEBA	10...	Power Supply Status 0x63
2020-05-19 12:07:43 AM	Active	High	Power Supply Failure	Host	NetworkHybrid	10...	Power Supply Status 0x62
2020-05-19 12:07:03 AM	Active	High	Power Supply Failure	Host	LogHybrid	10...	Power Supply Status 0x63
2020-05-18 11:41:13 PM	Active	High	Power Supply Failure	Host	ESAPrimary	10...	Power Supply Status 0x63
2020-05-18 10:58:44 AM	Active	High	Power Supply Failure	Host	SA	10...	SystemInfo/Power Supply Status 0x86
2020-06-04 12:01:47 AM	Cleared	Critical	Reporting Engine Service Stopped	Reporting Engine	1155A	10...	ProcessInfo/Service Status
2020-06-03 10:59:07 PM	Cleared	Critical	Host Unreachable	Host	SSCoreWithDAC	10...	ProcessInfo/Appliance Down
2020-06-03 04:12:07 PM	Cleared	Critical	Host Unreachable	Host	WindowsLC	10...	ProcessInfo/Appliance Down
2020-06-02 01:44:23 PM	Cleared	Critical	ContextHub Service Stopped	ContextHub Server	ESA	10...	ProcessInfo/Service Status
2020-06-02 10:29:27 AM	Cleared	Critical	Correlation Server Stopped	ESA Correlation	ESA	10...	ProcessInfo/Service Status

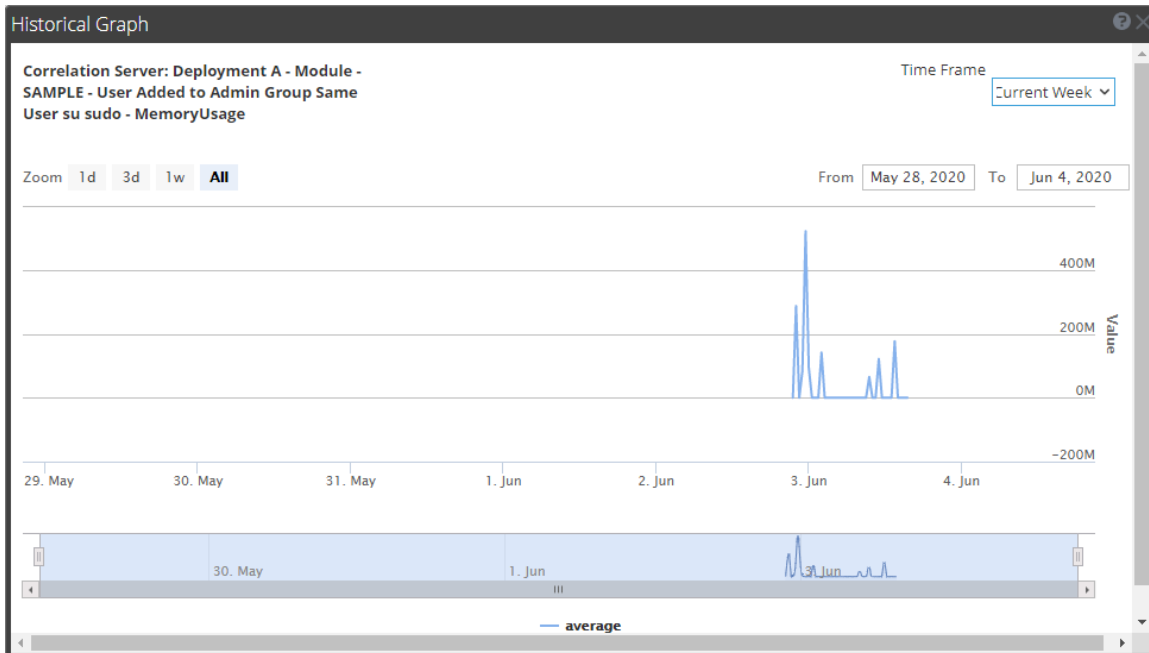
- Go to  (Admin) > **Health & Wellness** > **System Stats Browser** to see the memory metrics for each rule's performance. To view the metrics, enter the following and click **Apply**:

Host	Component	Category
<your host>	Correlation Server	Correlation Engine Metrics

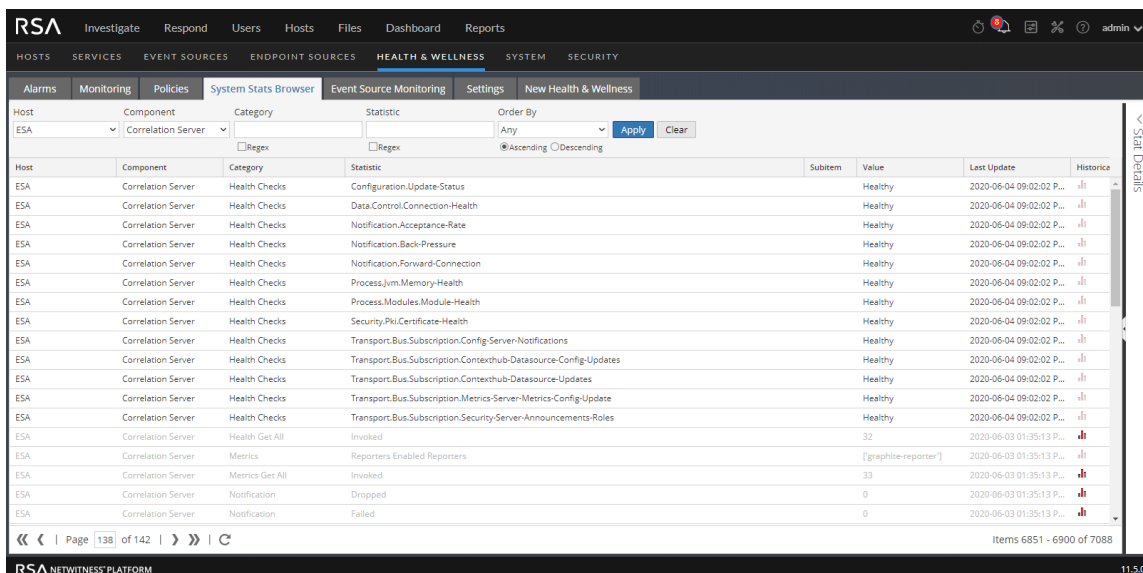
Host	Component	Category	Statistic	Subitem	Value	Last Update	Historical Graph
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Logons from Same Source IP with Unique Usernames - StatementFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - Deployed	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - DisplayName	Multiple Failed Pr...	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - Enabled	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - LastTimeAlertFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - MemoryUsage	564	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - StatementFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - Deployed	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - DisplayName	Multiple Intrusio...	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - Enabled	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - LastTimeAlertFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - MemoryUsage	88	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - StatementFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - Deployed	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - DisplayName	Multiple Login Fai...	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - Enabled	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - LastTimeAlertFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - MemoryUsage	176	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - StatementFired	0	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures by Administrators to Domain Controller - Deployed	true	2019-02-14 06:59:11 PM		
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures by Administrators to Domain Controller - DisplayName	Multiple Login Fai...	2019-02-14 06:59:11 PM		

The name of the rule is in the **Statistic** column and the memory usage in bytes is in the **Value** column.

3. Click  to view a historical view of memory usage for the rule in the **Historical Graph** column.



4. In the **System Stats Browser**, you can also see details of your ESA Correlation service performance.



Select your host, and use the following filters to view the following statistics:

Host	Component	Category	Statistic	Example
<your host>	Host	SystemInfo	CPU Utilization	49.74%
<your host>	Host	SystemInfo	Memory Utilization	72.87%
<your host>	Host	SystemInfo	Used Memory	183.14 GB
<your host>	Host	SystemInfo	Total Memory	251.65 GB
<your host>	Host	SystemInfo	Uptime	289546, 3 days 8 hours 25 minutes 46 seconds
<your host>	Correlation Server	Process jvm	Memory Total Max	163 GB
<your host>	Correlation Server	Process jvm	Memory Total Used	13.50 GB
<your host>	Correlation Server	ProcessInfo	CPU Utilization	0.3%
<your host>	Correlation Server	ProcessInfo	Maximum Memory	251.65 GB
<your host>	Correlation Server	ProcessInfo	Memory Utilization	151.87 GB

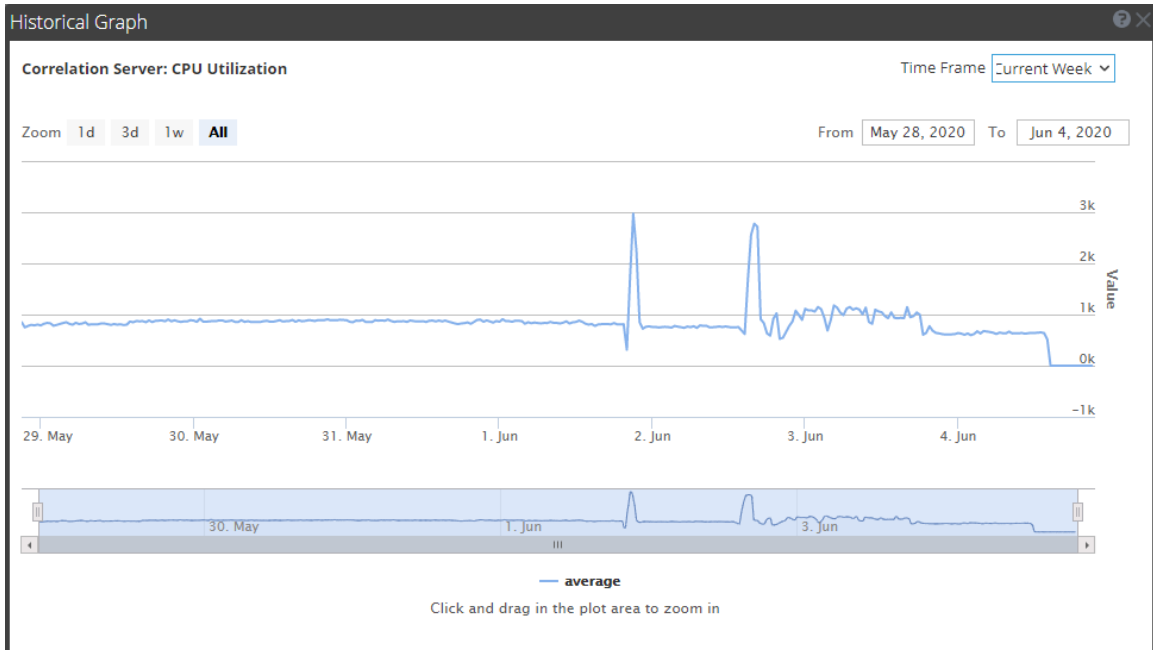
The following figure shows the location of the ESA Correlation service CPU and Memory Utilization statistics.

The screenshot shows the RSA NetWitness Platform interface. The 'System Stats Browser' tab is active, displaying a table of statistics for the 'Correlation Server' component. The table includes columns for Host, Component, Category, Statistic, Subitem, Value, Last Update, and Historical Graph. The 'CPU Utilization' and 'Memory Utilization' statistics are highlighted with red boxes. The 'CPU Utilization' value is 0.3% and the 'Memory Utilization' value is 151.87 GB. Both have 'Historical Graph' icons (a small bar chart) next to their last update timestamps.


Host	Component	Category	Statistic	Subitem	Value	Last Update	Historical Graph
ESA	Correlation Server	ProcessInfo	Build Date		2020-Mar-31 23:22:20	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	CPU Utilization		0.3%	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	Maximum Memory		251.65 GB	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	Memory Utilization		151.87 GB	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	Overall Processing Status Indicator		WORKING	2020-06-04 09:20:31 P...	
ESA	Correlation Server	ProcessInfo	Overall Service Status Indicator		WORKING	2020-06-04 09:20:31 P...	
ESA	Correlation Server	ProcessInfo	Running Since		2020-Jun-04 03:53:24	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	Service Status		started	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	Service Version		11.5.0.0	2020-06-04 09:20:32 P...	
ESA	Correlation Server	ProcessInfo	Uptime		62827, 17 hours 27 minutes 7 seconds	2020-06-04 09:20:32 P...	

- Click to view a historical view of CPU and memory utilization.

The following figure shows the historical graph of **CPU utilization**.





The following figure shows the historical graph of **Memory Utilization**.



If you are having a problem with memory or CPU utilization, continue to step 3.

Step 3: Bring up your ESA Services

1. Go to  (Admin) > Services, select your ESA service, and then select  > Start.
2. Return to the ESA Service to troubleshoot which rules have created memory issues.

If your ESA service is stopping and restarting in a loop, you may need to call Customer Support to get the services to start.

If you are able to start your ESA service without a shutdown, continue to step 4.

Step 4: Check the Alerts and Events Volume

After you are able to restart your ESA service without an immediate shutdown, you can review the stats for your rules to see which rules are consuming too many resources. Sometimes, ESA services fail because a rule is generating too many alerts or a rule is matching too many events. Check for both of these issues if you have determined that memory usage is causing your ESA service to shut down.

View Alert Summaries

Rules that generate a high volume of alerts can overwhelm the system and cause it to fail or restart. To view the alert summaries, go to **Respond > Alerts**. In the **Filters** panel on the left, in the **Alert Names** section, select the alert name for the rule. The number of alerts with that name appears at the bottom of the Alerts list results. If the number is significantly high for a particular rule, you need to disable the rule and rewrite it to be more efficient.

The screenshot shows the RSA NetWitness Alerts interface. The 'ALERTS' tab is active, and a filter is applied for 'Direct Login To an Administrative Account'. The filter is highlighted with a red box. The table below shows one alert:


CREATED	SEVERITY	NAME	SOURCE	# EVENTS	HOST SUMMARY	INCIDENT ID
06/02/2020 05:00:57 am	100	Direct Login To an Administrative Account	Event Stream Analysis	1	computer	INC-8

At the bottom of the filter panel, there are buttons for 'Reset', 'Save', and 'Save as...'. A status bar at the bottom right indicates 'Showing 1 out of 1 items | 0 selected'.

To clear your filter, click **Reset**.

View Events Matched

Sometimes a rule matches too many events, which can use up excessive memory. This typically occurs if you create a large event window where a great number of events accumulate without triggering an alert. This is a problem because each event is stored in memory while the rule waits for the alert to trigger. To

check for this issue, go to  **(Configure) > ESA Rules > Services**. From there, you can see the number of events that were matched in the **Events Matched** column for the deployment. If a high number of events were matched for a given rule, you can investigate the rule further to see if you can make it more efficient.

The screenshot shows the RSA NetWitness Services page for the 'esaSecondary - ESA Correlation' rule. The 'Deployment A' tab is active, showing the following statistics:

Engine Stats	Rule Stats	Alert Stats
Esper Version: 8.4.0	Rules Enabled: 75	Notifications: 0
Time: 2020-05-05T19:44:43	Rules Disabled: 0	Message Bus: 1
Events Offered: 74750	Events Matched: 29	
Offered Rate: 0 per second / 8,557 max		
Status: Active		

Below the statistics is a table of 'Deployed Rule Stats':


Enabled	Name	Rule Type	Trial Rule	Last Detected	Events Matched	Memory Usage	CPU %
<input type="checkbox"/>	AWS Critical VM Modified	Esper	Yes		0	0 bytes	0.017
<input type="checkbox"/>	AWS Permissions Modified Followed By Instan...	Esper	Yes		0	168 bytes	0.069
<input type="checkbox"/>	Account Added to Administrators Group and ...	Esper	Yes		0	168 bytes	0.134
<input type="checkbox"/>	Account Removals From Protected Groups on ...	Esper	Yes		0	64 bytes	0.047
<input type="checkbox"/>	Aggressive Internal Database Scan	Esper	Yes		0	64 bytes	0.087
<input type="checkbox"/>	Aggressive Internal NetBIOS Scan	Esper	Yes		0	4.22 KB	0.851
<input type="checkbox"/>	Aggressive Internal Web Portal Scan	Esper	Yes		0	1016 bytes	0.253
<input type="checkbox"/>	BYOD Mobile Web Agents Detected	Esper	Yes		0	64 bytes	0.079
<input type="checkbox"/>	Backend Activity Detected	Esper	Yes		0	0 bytes	0.014
<input type="checkbox"/>	Cerber Ransomware	Esper	Yes		0	896 bytes	0.314

The page also includes navigation controls for the table, showing 'Page 1 of 1' and 'Page Size 100'. The bottom status bar indicates 'Displaying 1 - 75 of 75'.




Step 5: Disable and Repair the Rule that Caused Issues

Once you have determined the rules that need to be rewritten, disable them and rewrite rules so that they don't generate such a high volume of alerts or events. For pointers on how to write more efficient rules, see [Best Practices](#).


Disable Rules

1. To disable rules, go to  (**Configure**) > **ESA Rules** > **Services**, and select the rules you want to disable in the **Deployed Rules Stats** field.
2. Select **Disable** to disable the rules.

Edit Rules


1. To repair the rules, go to  (**Configure**) > **ESA Rules** > **Rules tab** > **Rule Library**.
2. For each rule that you repair, do the following:
 - a. Select the rule to edit and then select   > **Edit**.
 - b. Edit the rule to be more efficient. For instructions on creating rules, see [Add Rules to the Rule Library](#)
 - c. When you are satisfied with your rule, you can save the rule as a trial rule to ensure that any memory issues do not affect ESA services performance. To do this, follow the steps listed in [Working with Trial Rules](#).

Deploy Rules

1. Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
2. In the options panel on the left, select the deployment that contains the rule.
3. In the Deployment view, the rule that you changed shows a status of Updated. Click **Deploy Now**. The rule status changes to Deployed.

Verify that the Rules are Enabled

After you deploy the ESA rules, they should automatically show as enabled. If not, you can enable the rules.

1. Go to  (**Configure**) > **ESA Rules** > **Services** tab, and select the ESA service in the options panel.
2. On the deployment tab for the deployment that contains the rules, in the Deployed Rule Stats section, look at the status of the rules in the Enable column. Enabled rules show a green circle. If the rules show a white circle, you can enable the rules.
3. To enable rules, select the rules you want to enable and select **Enable** above the table.

(Optional) Check the ESA Correlation Log Files for More Information

Once you verify that your services are down and some potential causes for the system going down, check to see if the service is stopping and restarting in a loop. To do this, go to the ESA Correlation logs. You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`.

ESA Rule Troubleshooting with Nw-Shell


Note: This procedure applies to NetWitness Platform 11.3 and later versions.


The ESA Correlation service replaces the Event Stream Analysis service in NetWitness version 11.3 and later. As a result of this change, some settings are no longer available in the user interface. In addition to the standard troubleshooting methods available, you can use the **nw-shell** utility to perform advanced troubleshooting of the ESA Correlation service and rules. For detailed information on the nw-shell utility, see the *RSA NetWitness Shell User Guide*.



- [Find Your Engine Name for Nw-Shell](#)
- [Connect to an ESA Correlation Server](#)
- [View the Contents of a Named Window](#)
- [See the Method Input and Output](#)

Find Your Engine Name for Nw-Shell

Follow these steps to find your engine name using your ESA rule deployment name. Your engine name is required for ESA Nw-Shell troubleshooting. Locate the names of each deployment that you plan to troubleshoot.

1. In the NetWitness UI, go to  (Admin) > **Health & Wellness** > **System Stats Browser**.
2. In the System Stats Browser, use the following filters and then click **Apply**.
 - a. **Host:** Select your ESA host.
 - b. **Component:** Select **Correlation Server**.
 - c. **Statistic:** Type **- Name** (put a space between the dash and name).
3. In the **Statistic** column, locate your deployment followed by **- Name**. The name in the **Value** field is YOUR ENGINE NAME.



Host	Component	Category	Statistic	Subitem	Value	Last Update	Historical Graph
ESA	Correlation Server	Correlation Engine Metrics	Deployment - Name		deployment-a-sa-managed	2020-07-13 06:38:42 PM	
ESA	Correlation Server	Correlation Engine Metrics	DeploymentA - Stream - Name		deployment-a-sa-managed-stream	2020-07-13 06:38:42 PM	

In the above example, the deployment name is DeploymentA and the engine name is deployment-a-sa-managed.

Connect to an ESA Correlation Server

1. Log in to nw-shell:
 - a. Connect via SSH to the NW server (head unit).
 - b. After logging in and getting the command prompt, type `nw-shell`.
2. Connect to ESA:
 - a. Go to your ESA physical host and type the command:
`cat /etc/netwitness/correlation-server/service-id`
 - b. Go back to the NW server and nw-shell and connect to the correlation server:
`connect --service correlation-server.ID`
3. Log in to ESA with the admin user credentials.
 - a. Type `login`.
 - b. Enter the username and password of the admin user.

```
[root@SAUII ~]# nw-shell
RSA
RSA Netwitness Shell. Version: 4.14.1-SNAPSHOT
offline » connect --service correlation-server.6a0f8e20-18da-45e7-acc3-90cc982e1cfb
INFO: Connected to correlation-server (6a0f8e20-18da-45e7-acc3-90cc982e1cfb)
correlation-server:Folder:/rsa » login
user: admin
password: *****
```

View the Contents of a Named Window

Use the `execute-query` method to view the contents of a named window.

1. After you are connected to the ESA correlation service and authenticated, type:
`cd /rsa/correlation/engine/execute-query`
2. Type `invoke '{"engineName":"<YOUR ENGINE NAME>", "query":"<YOUR QUERY>"}`
 - Where `<YOUR ENGINE NAME>` = the engine name that you located in [Find Your Engine Name for Nw-Shell](#).
 - Where `<YOUR QUERY>` = the select statement into the named window.

Example: `invoke '{"engineName":"esa-sa-managed", "query":"select * from UserLoginProfile"}'`

```
[root@SAUII ~]# nw-shell
RSA
RSA Netwitness Shell. Version: 4.14.1-SNAPSHOT

offline » connect --service correlation-server.6a0f8e20-18da-45e7-acc3-90cc982e1cfb
INFO: Connected to correlation-server (6a0f8e20-18da-45e7-acc3-90cc982e1cfb)
correlation-server:Folder:/rsa » login
user: admin
password: *****
admin@correlation-server:Folder:/rsa » cd /rsa/correlation/engine/execute-query
admin@correlation-server:Method:/rsa/correlation/engine/execute-query » invoke '{"engineName":"esa-sa-managed", "query":"select * from UserLoginProfile"'
[
  {
    "name" : "success",
    "cnt" : 247,
    "value" : "smithj"
  },
  {
    "name" : "success",
    "cnt" : 247,
    "value" : "doej"
  },
  {
    "name" : "failure",
    "cnt" : 247,
    "value" : "u408798"
  }
]
```

See the Method Input and Output

Type `show` at the command line to see the input and output expected. All method invocation must be prefaced by `invoke`.

```
correlation-server:Method:/rsa/correlation/engine/execute-query » show
```

Method	/rsa/correlation/engine/execute-query
output	java.util.List<java.util.Map<java.lang.String, java.lang.Object>>
input	com.rsa.netwitness.correlation.api.engine.QueryRequest
description	Execute the query in the given request. @param request (@link QueryRequest). @return (@link List) of Event (@link Map).

Metric	Value
invoked	7
timer	1454347.3481570843

Obtain Correlation Server Metrics for ESA Rule Deployment Troubleshooting Using Nw-Shell

Note: This procedure is available in NetWitness Platform version 11.4.1 and later.

You can use Nw-Shell to view ESA Correlation Server metrics for each of your ESA rule deployments. These metrics show the number of sessions behind for the deployment data sources as well as the memory usage for the rules in the deployment.

1. Find the engine name to use for Nw-Shell. See [Find Your Engine Name for Nw-Shell](#).
2. Connect to an ESA Correlation Server. See [Connect to an ESA Correlation Server](#).
3. After you are connected to the ESA correlation service and authenticated, type:
`cd /rsa/correlation/service/stats/get-condensed-metrics`
4. Type `invoke '<YOUR ENGINE NAME>'`
Where `<YOUR ENGINE NAME>` is the engine name that you located in [Find Your Engine Name for Nw-Shell](#).

Here is an example of the metrics output that you can obtain for your ESA rule deployment:

```
{
  "engineName" : "esa",
  "eventsOffered" : 1650,
  "maxEventsRate" : 2.337019271237945,
  "eventsRate" : {
    "count" : 1650,
    "oneSecRate" : 0.02129597415982235,
    "meanRate" : 0.05795554387594279,
    "oneMinuteRate" : 0.15485195471608212,
    "fiveMinuteRate" : 0.12419048320215775,
    "fifteenMinuteRate" : 0.11923922260543295
  },
  "streamMetrics" : {
    "pollingRate" : {
      "count" : 30100,
      "meanRate" : 1.0572440914676082,
      "oneMinuteRate" : 1.1497867142612552,
      "fiveMinuteRate" : 1.1213207743063618,
      "fifteenMinuteRate" : 1.1178158863433476
    },
    "positionTracking" : 475,
    "polling" : 30100,
    "bufferedRecords" : 0,
    "incomingRecords" : {
      "count" : 1650,
      "meanRate" : 0.05796815004652209,
      "oneMinuteRate" : 0.16275810590703457,
      "fiveMinuteRate" : 0.1253818211054517,
      "fifteenMinuteRate" : 0.11956773913684943
    },
    "outgoingRecords" : {
      "count" : 1650,
      "meanRate" : 0.05796815072883537,
      "oneMinuteRate" : 0.16275810590703457,
      "fiveMinuteRate" : 0.1253818211054517,
      "fifteenMinuteRate" : 0.11956773913684943
    },
    "sourceMetrics" : {
      "nw://admin@10.10.10.01:50005?compression=0&compressionLevel=6" : {
        "bufferedRecords" : 0,
        "incomingRecords" : {
          "count" : 1650,
          "meanRate" : 0.05796819356968545,
          "oneMinuteRate" : 0.16275810590703457,
          "fiveMinuteRate" : 0.12538181849661215,
          "fifteenMinuteRate" : 0.11956716772097528
        }
      }
    }
  }
}
```

```
    },
    "outgoingRecords" : {
      "count" : 1650,
      "meanRate" : 0.05796819300115741,
      "oneMinuteRate" : 0.16275810590703457,
      "fiveMinuteRate" : 0.12538181849661215,
      "fifteenMinuteRate" : 0.11956716772097528
    },
    "sessionsBehind" : 2,
    "sessionLastId" : "1645",
    "sessionRate" : "0",
    "lastReceivedSessionId" : 1645
  }
}
},
"ruleMetrics" : [
  {
    "ruleName" : "create_persist",
    "memoryUsage" : 104,
    "cpuLockedTimePercentage" : 50.286,
    "cpuLockedTimeNanos" : 138253,
    "statementFired" : 0,
    "alertsFired" : 0
  },
  {
    "ruleName" : "test_persist",
    "memoryUsage" : 72,
    "cpuLockedTimePercentage" : 39.44,
    "cpuLockedTimeNanos" : 108433,
    "statementFired" : 0,
    "alertsFired" : 0
  },
  {
    "ruleName" : "test-per",
    "memoryUsage" : 0,
    "cpuLockedTimePercentage" : 10.275,
    "cpuLockedTimeNanos" : 28249,
    "statementFired" : 0,
    "alertsFired" : 0
  }
]
}
```

View Memory Metrics for Rules

This topic tells ESA rule writers how to view memory metrics for an ESA Correlation service and its associated ESA rules. You can see estimated memory usage for each rule running on a server, and you can use this information to modify your rule statements and conditions if they use too much memory.

Rules can sometimes consume more memory than you expect, causing ESA to slow down or stop. To see approximately how much memory a rule is using, you can view estimated memory usage for each rule in the Health & Wellness System Stats browser (you need permissions to access this module). You can use this information to modify your rules to be more efficient.

At a high level, you need to complete the following steps to use memory metrics to troubleshoot memory usage for rules:

1. Ensure you have the correct permissions to view the Health & Wellness module. For information on roles and permissions, see [ESA Permissions](#).
2. View the memory statistics in Health & Wellness.
3. (Recommended) Configure Health & Wellness ESA policies to send an email if memory thresholds are exceeded. See "Manage Policies" in the *System Maintenance Guide* for instructions on sending email notifications.
4. Use the memory metrics data to modify rules to be more efficient, if necessary.

Note: You can also view memory metrics for ESA rules in the  **(Configure) > ESA Rules > Services** tab. See [View Stats for an ESA Service](#).

Prerequisites

The following are requirements for using memory metrics:

- You must have the appropriate permissions to view Health & Wellness statistics.
- (Recommended) Configure the ESA Health & Wellness policy to send an email when memory thresholds are exceeded.

Note: Memory Metrics is always on for the ESA Correlation service; you do not have to enable it.

View Health Statistics and Trends for ESA Correlation in New Health & Wellness

In NetWitness version 11.5 and later, New Health & Wellness provides improved and intuitive dashboards, monitors, and visualizations. The ESA Correlation Overview dashboard provides health statistics and trends on ESA rule deployments.



For more information, see "Monitor New Health and Wellness" and "Appendix A: New Health and Wellness Dashboards / ESA Correlation Overview Dashboard" in the *System Maintenance Guide*.

View Memory Metrics for an ESA Correlation Service in Health & Wellness

1. Go to  (Admin) > Health & Wellness > Monitoring tab.

Host	Status	CPU	Memory
ELH	Ready	0.79%	73.11 GB/125.92 GB
ESA	Ready	1.83%	11.04 GB/125.92 GB
PH	Ready	3.2%	32.59 GB/125.92 GB
SA	Ready	3.73%	42.90 GB/125.92 GB

Service	Health Status	Rate	Name	Service Type	CPU	Memory Usage	Uptime
Ready	●	--	ESA - Contexthub Server	Contexthub Server	0.5%	5.00 GB	2 weeks 5 days 22 hours 20 minutes 47 seconds
Ready	●	--	ESA - ESA Correlation	ESA Correlation	23.9%	3.32 GB	2 weeks 5 days 22 hours 20 minutes 47 seconds
Ready	●	0	PH - Concentrator	Concentrator	49.9%	285.21 MB	2 weeks 5 days 22 hours 20 minutes 47 seconds
Ready	●	0	PH - Decoder	Decoder	0.9%	643.26 MB	2 weeks 5 days 22 hours 16 minutes 49 seconds

2. Locate your host and click the link in the **Name** field for your ESA Correlation service, for example, **ESA- ESA Correlation**.

Service	Health Status	Rate	Name	Service Type	CPU	Memory Usage	Uptime
Ready	●	--	ESAPrimary1 - Event Stream Analytics Server	Entity Behavior Analytics	0.5%	2.50 GB	2 days 17 hours 39 minutes 2 seconds
Ready	●	--	ESAPrimary1 - ESA Correlation	ESA Correlation	4.8%	2.88 GB	2 days 17 hours 49 minutes 14 seconds

- On the tab for your ESA host, click the **Health Stats** tab.
You can view the health status of the ESA Correlation service.

The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation bar has tabs for 'Alarms', 'Monitoring', 'Policies', 'System Stats Browser', 'Event Source Monitoring', 'Settings', 'New Health & Wellness', and 'ESA'. The 'Health Stats' tab is selected, showing the following details:

Service			
CPU	23.9%	Used Memory	3.32 GB
Running Since	2020-Jun-04 21:53:57	Max Process Memory	125.92 GB
Build Date	2020-May-26 22:22:01	Version Information	11.5.0.0

Health Stats			
Configuration Update Status	Healthy	Process Modules	Healthy
Process JVM Memory	Healthy	Security PKI Certificate	Healthy
Data Connection	Healthy		

- Click the **JVM** tab.
You can view the JVM total memory used by the selected ESA Correlation service.


The screenshot shows the RSA NetWitness Platform interface with the 'JVM' tab selected. The details are as follows:

Service			
CPU	24.9%	Used Memory	3.32 GB
Running Since	2020-Jun-04 21:53:57	Max Process Memory	125.92 GB
Build Date	2020-May-26 22:22:01	Version Information	11.5.0.0


JVM			
JVM Total Memory Max	81.00 GB	JVM Total Memory Used	1.15 GB

Note: You can also view memory metrics for the ESA Correlation service in the  **(Configure) > ESA Rules > Services** tab. See [View Stats for an ESA Service](#).











View Memory Metrics for an ESA Correlation Service and its ESA Rules

1. Go to  (Admin) > Health & Wellness > System Stats Browser.
2. To view memory metrics for an ESA Correlation service, in the **Host** field, select your ESA host. Select **Correlation Server** for **Component**, enter **ProcessInfo** for **Category**, and then click **Apply**.

Host	Component	Category
<your host>	Correlation Server	ProcessInfo



The screenshot shows the RSA System Stats Browser interface. The 'Host' is set to 'ESA', 'Component' is 'Correlation Server', and 'Category' is 'ProcessInfo'. The 'Statistic' field is set to 'Memory Utilization'. The table below displays various system statistics for the selected host and component.

Host	Component	Category	Statistic	Subitem	Value	Last Update	Historical Graph
ESA	Correlation Server	ProcessInfo	Build Date		2020-May-26 22:22:01	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	CPU Utilization		26.9%	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Maximum Memory		125.92 GB	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Memory Utilization		3.32 GB	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Overall Processing Status Indicator		WORKING	2020-06-21 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Overall Service Status Indicator		WORKING	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Running Since		2020-Jun-04 21:53:57	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Service Status		started	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Service Version		11.5.0.0	2020-06-24 08:19:44 PM	
ESA	Correlation Server	ProcessInfo	Uptime		1722347, 2 weeks 5 d...	2020-06-24 08:19:44 PM	

The **Memory Utilization** statistic shows the total memory in use by the ESA Correlation service.

- To view the historical memory usage for the ESA Correlation service, click the **Historical Graph** icon.




- To view the memory metrics for individual rules, in the **Category** field, enter **Correlation Engine Metrics** and click **Apply**.

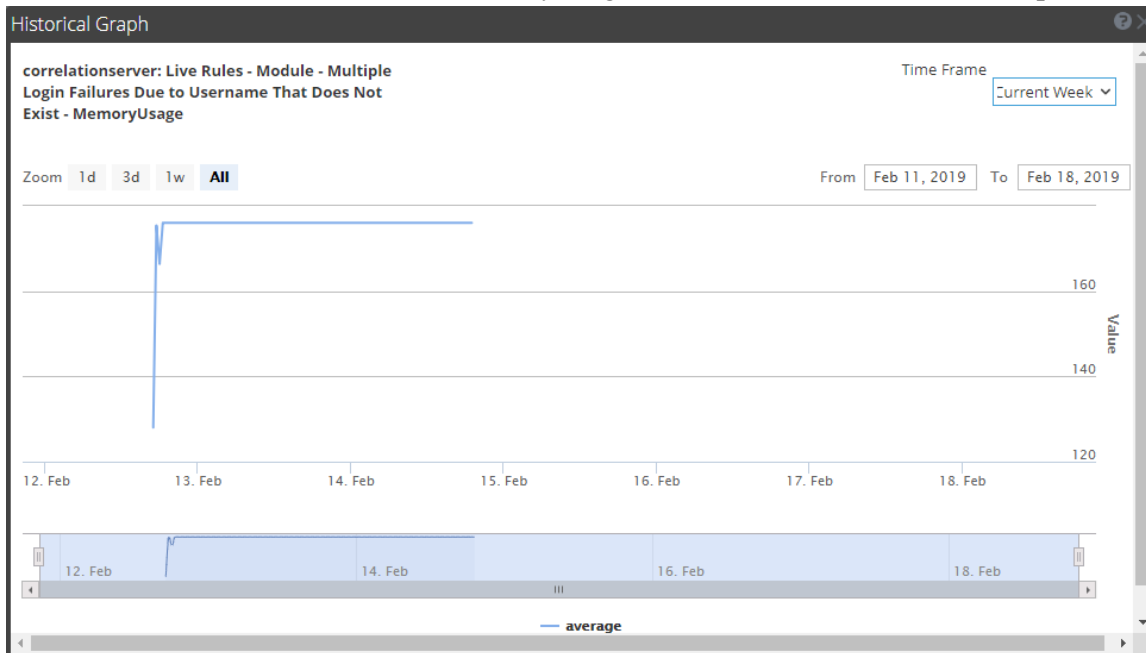
Host	Component	Category
<your host>	Correlation Server	Correlation Engine Metrics


Host	Component	Category	Statistic	Value	Last Update
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Logons from Same Source IP with Unique Usernames - StatementFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - Deployed	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - DisplayName	Multiple Failed Pr...	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - Enabled	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - LastTimeAlertFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - MemoryUsage	344	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Failed Privilege Escalations by Same User - StatementFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - Deployed	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - DisplayName	Multiple Intrusio...	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - Enabled	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - LastTimeAlertFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - MemoryUsage	88	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Intrusion Scan Events from Same User to Unique Destinations - StatementFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - Deployed	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - DisplayName	Multiple Login Fai...	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - Enabled	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - LastTimeAlertFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - MemoryUsage	176	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures Due to Username That Does Not Exist - StatementFired	0	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures by Administrators to Domain Controller - Deployed	true	2019-02-14 06:59:11 PM
ESA10333	Correlation Server	Correlation Engine Metrics	Live Rules - Module - Multiple Login Failures by Administrators to Domain Controller - DisplayName	Multiple Login Fai...	2019-02-14 06:59:11 PM

The name of the rule is in the **Statistic** column appended with **MemoryUsage** and the memory usage in bytes is in the **Value** column.

Note: The **Last Update** field reflects when Health & Wellness polls ESA. However, the Memory Metrics is not synchronized with the Health & Wellness polling. For example, if the memory threshold is exceeded on 2/10/19 at 12 p.m., but Health & Wellness polls at 2/10/19 at 12:10 p.m., the **Last Update** field will display a timestamp of 2/10/19 12:10 p.m.

5. Click  to view a historical view of memory usage for the rule in the **Historical Graph** column.



Note: You can also view memory metrics for ESA rules in the  **(Configure) > ESA Rules > Services** tab. See [View Stats for an ESA Service](#).

How ESA Handles Sensitive Data

This topic explains how ESA treats sensitive data, such as usernames or IP address, that it receives from Core services. The Data Privacy Officer (DPO) role can identify meta keys that contain sensitive data and should display obfuscated data. ESA does not display or store sensitive meta. Consequently, ESA will not pass sensitive data to NetWitness Respond.

Optionally, ESA can add an obfuscated version of the sensitive data to an event. For example, the DPO identifies `user_dst` as sensitive. ESA can add an obfuscated version, such as `user_dst_hash`, to an event. The obfuscated meta is not sensitive, so ESA will display and store it the same way as any other non-sensitive meta.

For more information on the strategy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

This topic explains the following:

- How ESA treats sensitive data it receives from Core services
- How to prevent sensitive data leaks in an Advanced EPL rule
- How to remove sensitive meta keys from global alerts

How ESA Treats Sensitive Data from Core Services

When ESA receives sensitive data from Core services, ESA passes on only the obfuscated version of the data. ESA does not store or show sensitive data.

The following features are impacted:

- Outputs – ESA does not forward sensitive data to outputs, which include alerts, notifications, and MongoDB storage.
- Advanced EPL rules – If an EPL statement creates an alias for a sensitive meta key, sensitive data will leak. This topic illustrates how this happens so you can avoid it.
- Enrichments – If a sensitive meta key is used in the join condition, sensitive data will leak. This topic illustrates how this happens so you can avoid it.

Advanced EPL Rule

If an EPL query statement renames a sensitive meta key, the data will not be protected.

ESA identifies a sensitive meta key by the name:

- `ip_src` is the sensitive meta key.
- `ip_src_hash` is the non-sensitive, obfuscated version.

To support data privacy, the sensitive meta key must not be renamed in an EPL query. If a sensitive meta key is renamed, the data will no longer be protected.

For example, in a rule such as `select ip_src as ip_alias...`, `ip_alias` contains the sensitive data but it is not protected because ESA only knows about `ip_src`, not `ip_alias`. In this case, IP addresses would not be obfuscated. Real values would be displayed.

Enrichment Source

When a sensitive meta key is used in a join condition, sensitive data can be displayed.

The enrichment database, which is the other part of the join condition, has one column that matches the sensitive meta key. This cross reference is to actual values not obscured values. Consequently, actual values are displayed.

In the following example, both parts of the join condition are highlighted.

Enrichments		+	⌵	-
Type	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name	
<input type="checkbox"/> GeolP	Default GeolP	ip_src	ipv4	



- `ip_src` contains sensitive data.
- `ipv4` will be added to the alert and exposed as non-sensitive data.

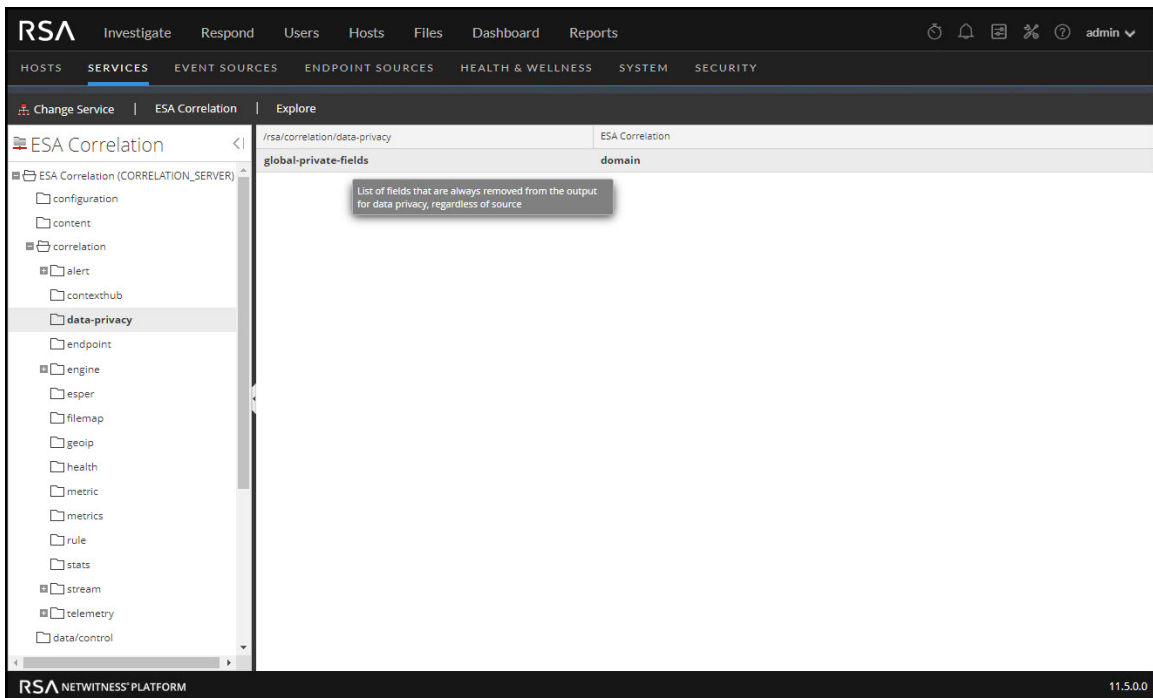
Because the `ipv4` value is the same as the `ip_src` value, `ipv4` contains and displays sensitive data.

How to Remove Sensitive Meta Keys Globally from All Alerts

Note: This procedure applies only to ESA Correlation Rules in NetWitness Platform 11.4 and later versions.

For data privacy reasons, it may be necessary to remove some sensitive meta keys from the alert output globally, regardless of the data source. In the ESA Correlation service, you can set the `global-private-fields` parameter to remove the meta keys from all alert output.

1. Go to  (Admin) > Services, and in the Services view, select an ESA Correlation service and then select  > View > Explore.
2. In the Explore view node list for the ESA Correlation service, select **correlation** > **data-privacy**.
3. In the `global-private-fields` parameter, add the sensitive meta keys that you want removed from all alerts.



The changes are effective immediately.

For more information on the strategy and benefits of obfuscating data, see the *Data Privacy Management Guide*.

ESA Rule Types

This topic describes each type of ESA rule, when to use them and the permissions each role has with them. The following table lists each type, describes it, and explains when to use it.

Rule Type	Description	When to Use
RSA Live ESA	RSA Live has a catalog of ESA rules that you can download and modify to run in your network.	Download RSA Live ESA rules to leverage rules that are already built. Modify the configurable parameters to customize to meet your requirements.
Rule Builder	In the rule builder, you define rule criteria in an easy-to-use interface.	Use the rule builder to create your first rules. You choose many of the rule conditions from lists.
Advanced EPL	With the Event Processing Language (EPL), you define rule criteria by writing a query.	Use advanced EPL rules to define rule criteria in the EPL syntax.
Endpoint Rule Bundle	An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. The rules in this bundle only apply to NetWitness Endpoint.	If you have NetWitness Endpoint, you can configure risk scoring to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see "Deploy Endpoint Risk Scoring Rules on ESA" in the <i>ESA Configuration Guide</i> . For complete information on configuring NetWitness Endpoint, see the <i>NetWitness Endpoint Configuration Guide</i> .

Sample Rules

Sample Rule Builder rules come with NetWitness and appear in the Rule Library. Use sample rules to get comfortable working with rules before creating your own. You can safely edit and deploy these sample rules.

Endpoint Risk Scoring Rules Bundle

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. These rules appear in the Rule Library with the sample rules. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

Trial Rules Mode

For any type of rule, you can select the Trial Rule setting as an additional safeguard. Trial rules get disabled if they exceed a memory threshold set by the administrator. Run a rule in trial mode to monitor memory usage and to disable the rule automatically if it uses more memory than the threshold allows.

The following figure shows the Trial Rule setting in the Rule Builder.

The screenshot shows the RSA NetWitness Platform Rule Builder interface. The 'Configure' tab is active, and the rule being edited is 'SAMPLE - Whitelist - From ...'. The 'Trial Rule' checkbox is checked and highlighted with a red box. The rule description states: 'Whitelist Germany from P2P software as detected by an intrusion detection device (IDS), intrusion prevention device (IPS), firewall or vulnerability scanner. This is the same as the rule "SAMPLE - P2P Software as Detected by an Intrusion Detection Device" with the addition of a separate whitelist condition which ignores attempts when the source IP appears to be from Germany.'

The 'Conditions' section contains the following table:

Statement	Occurs	Connector	Correlation Type	Meta	Meta
<input type="checkbox"/> Intrusion Log Message	1	AND			
<input type="checkbox"/> P2P Detection	1	AND			
<input type="checkbox"/> Whitelist Germany	1				

The 'Notifications' section is empty, showing 'No parameters to edit.'

The 'Enrichments' section contains the following table:

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
<input type="checkbox"/> GeoIP	Default: GeoIP	ip_src	ipv4

At the bottom of the interface, there are buttons for 'Save', 'Close', and 'Show Syntax', along with a note that '** required field'.

ESA Permissions

This topic lists all ESA permissions and shows which permissions are assigned to each pre-configured NetWitness role. User access is restricted based on roles and permissions assigned to roles.

- Administrators
- Operators
- Analyst
- Security Operations Center (SOC) Managers
- Malware Analysts (MA)
- Data Privacy Officer

There are four permissions for ESA:

- **Access Alerting Module:** Is required for any permission
- **View Rules:** Allows view-only permission for rules in the Rule Library
- **View Alerts:** Allows view-only permission for alerts ESA generates
- **Manage Rules:** Allows you to view, create, edit, and delete rules

The following table lists permissions for ESA and the roles to which they are assigned. Use this table to see how each role can work with rules and alerts.


Permission	Administrators	Operators	Analysts	SOC Mgrs	MA's	DPO's
Access Alerting Module	Yes	Yes	Yes	Yes		Yes
View Rules	Yes	Yes		Yes		Yes
View Alerts	Yes		Yes	Yes		Yes
Manage Rules	Yes	Yes		Yes		Yes

For more information on roles and permissions, see the *System Security and User Management Guide*.

Practice with Sample Rules

NetWitness comes with sample rules so analysts can become familiar with how rules look before they create their own rules. Use the sample rules to become familiar with the Rule Builder and to practice editing and deploying a rule.

Sample rules are installed in the Rule Library, which contains every rule you download or create. The following figure shows sample rules in the Rule Library.



These are the available sample rules:

- SAMPLE - Blacklist - From inside countries that are not the US, Non SMTP Traffic on TCP Port 25 Containing Executable
- SAMPLE - Non SMTP Traffic on TCP Port 25 Containing Executable
- SAMPLE - P2P Software as Detected by an Intrusion Detection Device
- SAMPLE - User Added to Admin Group Same User su Sudo
- SAMPLE - Whitelist - From outside of Germany, P2P Software as Detected by an Intrusion Detection Device.

Each name begins with SAMPLE to distinguish the rules that are installed with NetWitness from the rules you download and create.

Rule Library



The Rule Library shows the following information for a rule:

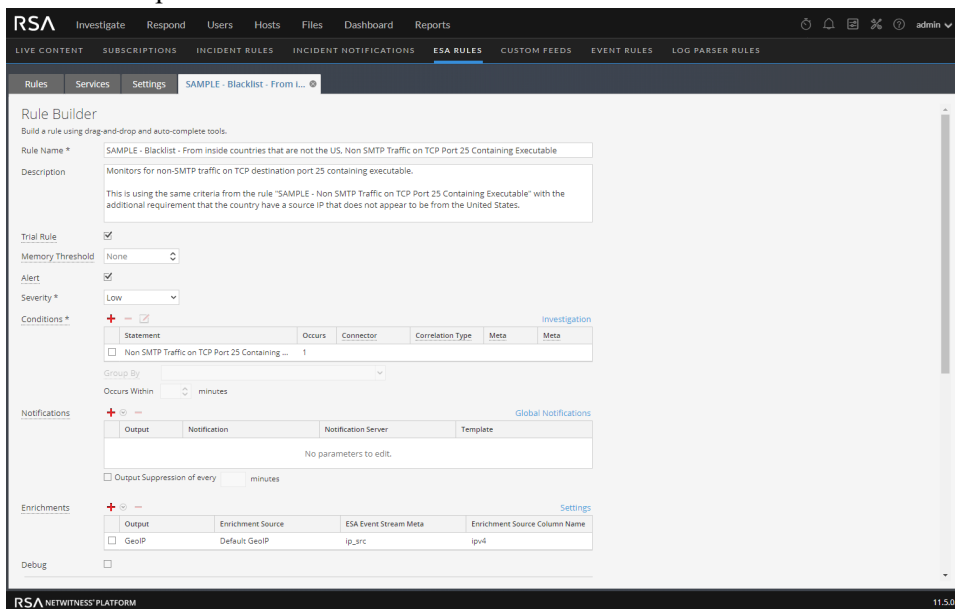
- **Name** summarizes the data or events the rule collects.
- **Description** explains the rule in more detail, although only the beginning shows in the Rule Library.
- **Trial Rule** indicates if trial mode is enabled or disabled for the rule.
- **Type** shows the origin of the rule, built in Rule Builder or Advanced EPL, downloaded from RSA Live, or Endpoint Rule Bundle.

The screenshot displays the RSA NetWitness Platform interface, specifically the Rule Library section. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation menu shows 'LIVE CONTENT', 'SUBSCRIPTIONS', 'INCIDENT RULES', 'INCIDENT NOTIFICATIONS', 'ESA RULES', 'CUSTOM FEEDS', 'EVENT RULES', and 'LOG PARSER RULES'. The 'ESA RULES' tab is active, and the 'Rule Library' sub-tab is selected. The interface shows a list of rules with the following columns: Rule Name, Description, Trial Rule, Type, and Actions. The selected rule is 'SAMPLE - Blacklist - From inside countries that are not the US, N...' with a description 'Monitors for non-SMTP traffic on TCP destination port 25 contai...'. The interface also includes a search filter, a 'GET RULES FROM RSA LIVE' button, and pagination controls at the bottom.

Rule Name	Description	Trial Rule	Type	Actions
<input type="checkbox"/> Endpoint Risk Scoring Rule Bundle	This bundle contains many rules which are used for risk scoring of E...	No	Endpoint Rule B...	
<input checked="" type="checkbox"/> SAMPLE - Blacklist - From inside countries that are not the US, N...	Monitors for non-SMTP traffic on TCP destination port 25 contai...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - Non SMTP Traffic on TCP Port 25 Containing Executable	Monitors for non-SMTP traffic on TCP destination port 25 containing ...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - P2P Software as Detected by an Intrusion Detection Device	P2P software as detected by an intrusion detection device (IDS) intru...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - User Added to Admin Group Same User su sudo	Alert when user is upgraded to one of admin groups (custom list of g...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - Whitelist - From outside of Germany, P2P Software as Dete...	Whitelist Germany from P2P software as detected by an intrusion de...	Yes	Rule Builder	

Practice with Sample Rules

1. Go to  **(Configure) > ESA Rules**.
The ESA Rules view is displayed with the Rules tab open.
2. In the **Rule Library**, double-click a sample rule or select a sample rule and click .
The rule is opened in Rule Builder.



3. To practice with a sample rule, refer to the following topics for detailed descriptions and procedures:
 - To familiarize yourself with the Rule Builder user interface, see [Rule Builder Tab](#) for a description of each field.
 - To learn how to edit a rule, see [Add a Rule Builder Rule](#) for a step-by-step procedure.
 - To deploy sample rule, see [Deploy Rules to Run on ESA](#) to learn how to associate the rule with an ESA service.

After you practice with sample rules, you will be able to download, create, and deploy your own rules.

Working with Trial Rules

The ESA Correlation service is capable of processing large volumes of disparate event data from Concentrators. However, when working with ESA Correlation rules, it is possible to create rules that use excessive memory. This can slow your ESA service or even cause it to shut down unexpectedly. To ensure that rules do not use excessive memory, you can enable them as trial rules. You should disable the trial rule setting only after testing the new rule in your environment during times of both normal and peak network traffic.

You can set a global threshold of the percentage of memory that trial rules may use. If that configured memory threshold is exceeded, all trial rules are disabled automatically. To configure the memory threshold, see "Change Memory Threshold for All Trial Rules" in the *ESA Configuration Guide*.

For suggestions on creating more efficient rules, see "Best Practices for Writing Rules" in [Best Practices](#).


By default, new rules and RSA Live rules that you import are configured as trial rules. As a best practice, when you edit an existing rule, select the Trial Rule option, which allows you to deploy the rule with an added safeguard.


Note: Run a rule as a trial rule long enough to assess the performance during normal and peak network traffic.

Deploy Rules as Trial Rules

This topic explains to administrators how to enable trial rules when creating new rules or editing rules. Trial rules are automatically disabled if a specified total JVM memory utilization threshold is exceeded.

In NetWitness 11.4 and later, ESA trial rules no longer change status after an upgrade or deployment.

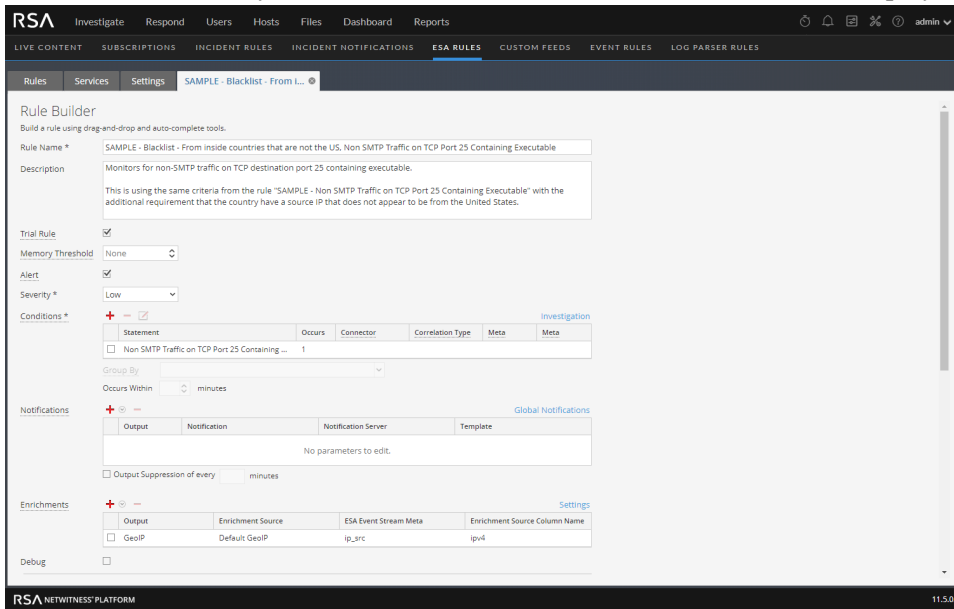
For example, if you change the status of a trial rule to disabled ( (Configure) > ESA Rules >

Services tab) and redeploy the ESA rule deployment ( (Configure) > ESA Rules > Rules tab), the trial rule remains disabled.

1. Go to  (Configure) > ESA Rules.


The Configure ESA Rules view is displayed with the Rules tab open.

- From the Rule Library, choose to add or edit a rule. The rule builder is displayed in a new tab.



- To make a new or existing rule a trial rule, select the **Trial Rule** checkbox.
- Add the rule conditions or modify the rule as needed. For instructions on editing rules, see [Add Rules to the Rule Library](#).
- Click **Save**.
- Ensure that trial rules are enabled for your ESA and that you are satisfied with the thresholds configured for trial rules.
The memory threshold is set in the configuration file. To configure it, see "Change Memory Threshold for All Trial Rules" in the *ESA Configuration Guide*.
 - The threshold is configured per ESA and is a percentage of Java Virtual Memory.
 - The configuration parameter, `fatal-percentage`, has a default value of 90.
- Optionally, you can set up the policies in Health and Wellness to send you an email notification if the total JVM memory utilization threshold is exceeded.

The next time you deploy the rule, it runs in trial rule mode.

Note: If a trial rule is disabled, you will need to go to the  **(Configure)** > **ESA Rules** > **Services** tab to re-enable the trial rules. For more instructions on re-enabling trial rules on a service, see [View ESA Stats and Alerts](#).

Add Rules to the Rule Library

This topic explains how to add each type of rule to the rule library. You must add a rule to the Rule Library before you can deploy it. Permission to manage rules is required for all tasks in this section. To add rules, you can download them from ESA Live, create a rule via the Rule Builder, or write advanced EPL rules.

For more details on each of these procedures, see:

- [Download Configurable RSA Live ESA Rules](#)
- [Add a Rule Builder Rule](#)
- [Add an Advanced EPL Rule](#)

In addition to deploying a rule, you can edit, duplicate, import, export, and remove a rule in the Rule Library. For details on these procedures, see [Working with Rules](#)

Download Configurable RSA Live ESA Rules

This topic explains how to download configurable rules from the NetWitness Live Content Management System so you can customize them to meet your needs.

RSA Live contains a catalog of rules. Each rule has configurable parameters so you can customize the rule for your environment. If RSA Live has a rule to detect events that you want to detect in your network, download the rule to save time. You can edit the configurable parameters and save the rule in your Rule Library. For detailed information about each rule, including whether the rule is for logs, packets, or both, see "RSA ESA Rules" at the following link:

<https://community.netwitness.com/t5/netwitness-platform-threat/rules/ta-p/677884>

This is an example of how each RSA Live ESA rule is described on RSA Live:

Rule Name	Description
Logins across Multiple Servers	<p>Detects logins from the same user across 3 or more separate servers within 5 minutes.</p> <p>The time window and number of unique destinations are configurable.</p>

As the name shows, the rule looks for logins across multiple servers. The description explains the rule criteria in more detail and specifies which parameters you modify.

Note: When a rule description includes a configurable parameter, the default setting for the parameter is used. In the example rule, the description states 5 minutes. However, the time window is configurable so 5 is the default number of minutes.


Prerequisites

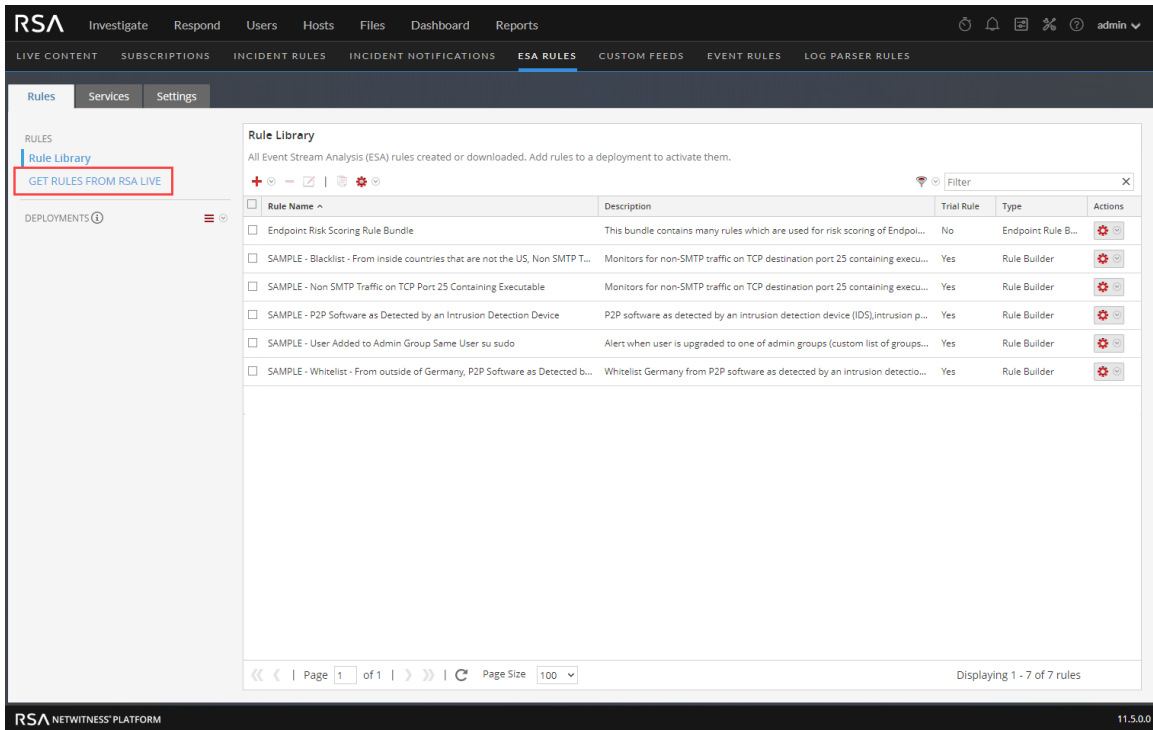
These are the prerequisites for downloading configurable RSA Live ESA rules;

- Have permission to manage rules
- Create a Live Account. See the *Live Services Management Guide* for details.
- Set up Live on NetWitness. See the *Live Services Management Guide* for details.
- Update your meta keys. See "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*.


Download RSA Live ESA Rules

Caution: Before you deploy the latest Live ESA rules, update your meta keys. See "Update Your ESA Rules for the Required Multi-Value and Single-Value Meta Keys" in the *ESA Configuration Guide*.

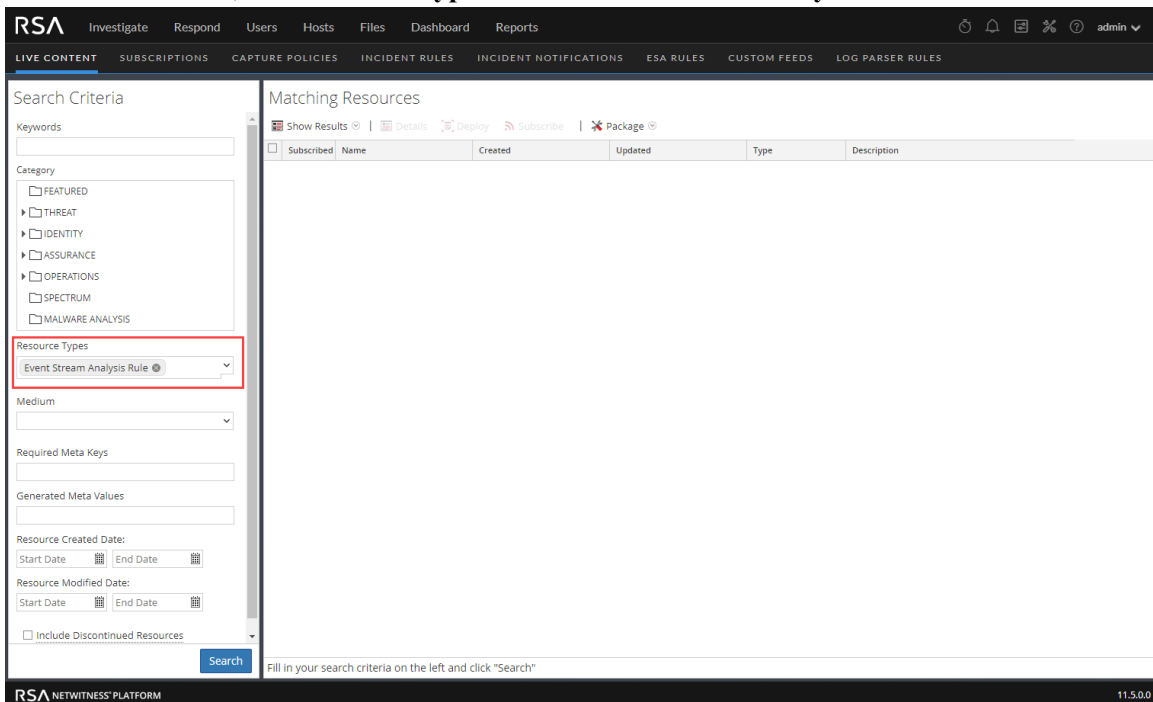
1. Go to  (Configure) > ESA Rules.
The Rules tab is displayed.



2. In the options panel, click **Get Rules from RSA Live**.

The Live Content Search view is displayed. (Alternatively, you can go to  **(Configure) > Live Content**.)

3. In **Search Criteria**, for **Resource Types** select **Event Stream Analysis Rule**.



- Specify any of the following criteria to find a rule to configure for your environment. For detailed information about each rule, including whether the rule is for logs, packets, or both, see "RSA ESA Rules" at the following link: <https://community.netwitness.com/t5/netwitness-platform-threat/rules/tag/677884>

For a detailed description of the search criteria, see "The Live Search View" in the *Live Services Management Guide*.

- Keywords
- Category
- Resource Types (Event Stream Analysis Rule)
- Medium (Log, Log and Packet, or Packet)
- Required Meta Keys
- Generated Meta Values
- Resource Created Date
- Resource Modified Date
- Include Discontinued Resources

- Click **Search**. Rules that match the search criteria are displayed in Matching Resources.

The screenshot displays the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. Below this, a secondary navigation bar lists various tool categories: 'LIVE CONTENT', 'SUBSCRIPTIONS', 'CAPTURE POLICIES', 'INCIDENT RULES', 'INCIDENT NOTIFICATIONS', 'ESA RULES', 'CUSTOM FEEDS', and 'LOG PARSER RULES'. The main interface is split into two panels. The left panel, titled 'Search Criteria', contains several search filters: 'Keywords' (with an input field), 'Category' (with a tree view showing 'FEATURED', 'THREAT', 'IDENTITY', 'ASSURANCE', 'OPERATIONS', 'SPECTRUM', and 'MALWARE ANALYSIS'), 'Resource Types' (set to 'Event Stream Analysis Rule'), 'Medium' (with a dropdown), 'Required Meta Keys' (with an input field), 'Generated Meta Values' (with an input field), 'Resource Created Date' (with 'Start Date' and 'End Date' pickers), 'Resource Modified Date' (with 'Start Date' and 'End Date' pickers), and a checkbox for 'Include Discontinued Resources'. A 'Search' button is at the bottom of this panel. The right panel, titled 'Matching Resources', shows a table of results. The table has columns for 'Subscribed', 'Name', 'Created', 'Updated', 'Type', and 'Description'. The 'RDP Inbound Traffic' rule is highlighted in red, indicating it is selected. Below the table, it indicates '102 Matching Resources'. The bottom of the interface shows the 'RSA NETWITNESS PLATFORM' logo and the version number '11.5.00'.

Subscribed	Name	Created	Updated	Type	Description
<input type="checkbox"/>	Advanced Rule Template	2013-12-24 11:23 AM	2013-12-24 11:23 AM	Event Stream Analysis Rule	This template is for advanced rule content module creation.
<input type="checkbox"/>	Basic Rule Template	2013-12-24 11:23 AM	2016-10-31 7:06 PM	Event Stream Analysis Rule	This template is for basic rule content module creation.
<input type="checkbox"/>	Internal Data Posting to 3rd...	2014-08-16 9:02 AM	2018-03-30 4:30 PM	Event Stream Analysis Rule	Detects when an internal IP address A receives an amount of data great...
<input type="checkbox"/>	AWS Permissions Modified ...	2015-08-12 10:57 PM	2019-07-22 8:50 PM	Event Stream Analysis Rule	Detects when an Amazon Web Services (AWS) permission is modified fol...
<input type="checkbox"/>	VM Clone After Multiple Ro...	2014-01-22 6:16 PM	2016-12-14 8:17 PM	Event Stream Analysis Rule	10.4 or higher. Alert if there are 3 root login failures to an ESX server foll...
<input type="checkbox"/>	No logs traffic from device ...	2014-02-27 11:23 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule	Detects when there is no traffic from a device for a specified time period...
<input type="checkbox"/>	Multiple Login Failures fro...	2014-03-14 10:44 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule	Detects when log events that contain multiple failed login events from th...
<input type="checkbox"/>	Malicious Account Creation...	2014-02-27 11:23 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule	Detects when a new account is created on a system and three authentic...
<input checked="" type="checkbox"/>	RDP Inbound Traffic	2014-02-27 11:24 AM	2017-11-02 1:38 PM	Event Stream Analysis Rule	Identifies RDP inbound traffic from one or more source IPs to 2 unit...
<input type="checkbox"/>	RDP traffic from Same sour...	2014-02-27 11:24 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule	Detects RDP traffic from the same source to multiple different destinac...
<input type="checkbox"/>	Suspicious Account Removal	2014-08-16 9:02 AM	2016-12-14 8:19 PM	Event Stream Analysis Rule	Detects a user account that has been added to an administrative group...
<input type="checkbox"/>	Malware Dropper	2015-08-04 8:13 AM	2017-12-18 9:42 PM	Event Stream Analysis Rule	This rule triggers upon download of pdf, java, rtf, or Microsoft Office file...
<input type="checkbox"/>	Windows Suspicious Admin...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	Detects when a user account is created, added to the Administrators grc...
<input type="checkbox"/>	Windows Suspicious Admin...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	Detects when a Windows user account is created, a shared object is acco...
<input type="checkbox"/>	User Account Created and ...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	Detects when a user account is created and then gets deleted within the...
<input type="checkbox"/>	krbtgt Account Modified on...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	Detects modification to the krbtgt account on domain controller. There...
<input type="checkbox"/>	Multiple Failed Logins from...	2018-03-27 2:01 PM	2018-03-27 3:40 PM	Event Stream Analysis Rule	Multiple failed logins from the same user, originating from multiple diffe...
<input type="checkbox"/>	Multiple Login Failures by A...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	This rule is triggered when a user enters Administrator credentials to log...
<input type="checkbox"/>	Multiple Failed Privilege Esc...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	Triggers after a user account fails privilege escalation multiple times with...
<input type="checkbox"/>	Multiple Login Failures by G...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	This rule is triggered when a user enters Guest credentials to log in to a...
<input type="checkbox"/>	Failed Logins Outside Busin...	2017-11-30 10:39 PM	2017-11-30 10:39 PM	Event Stream Analysis Rule	This rule is triggered when a user logs into a system after business hour...
<input type="checkbox"/>	Logins across multiple serv...	2018-03-27 2:00 PM	2018-03-27 2:00 PM	Event Stream Analysis Rule	Detects logins from the same user across multiple separate servers or h...
<input type="checkbox"/>	Privilege User Account Pass...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	Detects a logged modification of an administrative account password. TI...
<input type="checkbox"/>	Discrete Login By A Member...	2018-03-27 2:01 PM	2018-03-27 2:01 PM	Event Stream Analysis Rule	A successful interactive or remote interactive login to a user account...

6. Select each rule to download and click **Deploy**.

Matching Resources

Show Results | Details **Deploy** | Subscribe | Package

<input type="checkbox"/>	Subscribed	Name	Created	Updated	Type
<input type="checkbox"/>	no	No logs traffic from device i...	2014-02-27 11:23 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule
<input type="checkbox"/>	no	Multiple Login Failures fro...	2014-03-14 10:44 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule
<input type="checkbox"/>	no	Malicious Account Creation...	2014-02-27 11:23 AM	2016-12-14 8:18 PM	Event Stream Analysis Rule
<input checked="" type="checkbox"/>	no	RDP Inbound Traffic	2014-02-27 11:24 AM	2017-11-02 1:38 PM	Event Stream Analysis Rule

The Deployment Wizard is displayed.

Deployment Wizard

Resources Services Review Deploy

Total resources : 4

Resource Names	Resource Type	Dependency of
traffic_flow	Lua Parser	RDP Inbound Traffic
RDP_lua	Lua Parser	RDP Inbound Traffic
RDP Inbound Traffic	Event Stream Analysi...	
Investigation	Feed	RDP_lua,traffic_flow

Cancel Next

7. Follow the steps in the wizard. If you need more information, see "Deploy Resources in Live" in the *Live Services Management Guide*.


When you finish the steps in the wizard, the selected rules are displayed in the Rule Library.

Customize an RSA Live ESA Rule

This topic explains how to configure parameters in an RSA Live ESA rule. When you download an RSA Live ESA rule, the rule appears in the Rule Library which includes the following columns:

- Rule Name
- Description
- Trial Rule

- Type
- Actions



Rule Library
All Event Stream Analysis (ESA) rules created or downloaded. Add rules to a deployment to activate them.



Rule Name	Description	Trial Rule	Type ^	Actions
<input type="checkbox"/> User Account Created and Deleted within an Hour	User account created and deleted within an hour	No	RSA Live ESA Rule	
<input type="checkbox"/> Port Scan Horizontal Log	Series of log events indicating a port scan.	Yes	RSA Live ESA Rule	

The type is RSA Live ESA Rule.

Prerequisites

- Administrator, Operator, SOC Manager, or DPO role permissions are required.
- Rules must be downloaded to the Rule Library.

Configure Parameters for an RSA Live ESA Rule

1. Go to  (Configure) > **ESA Rules** > **Rules** tab.
2. In the **Rule Library**, double-click an RSA Live ESA Rule or select the rule and click . The RSA Live ESA Rule tab is displayed.
3. (Optional) Change the following fields:
 - Rule Name
 - Description
 - Trial Rule (Enabled by default. RSA recommends you run a rule as a trial rule long enough to assess the performance during normal and peak network traffic.)
 - Alert (This option applies to 11.3 and later.) Select Alert to send an alert to Respond. Clear the checkbox if you do not want to send an alert to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.
 - Severity
 - Notifications
 - Enrichments
4. To configure the rule for your environment, in the **Parameters** section replace the default in the **Value** Column.

Name ^	Value
With this number of events	200
Within this number of seconds	60

5. Click **Save**.

Add a Rule Builder Rule

Each ESA rule is designed to detect something in your network and to generate an alert for it:

- User activity that is not allowed, such as attempting to download software that is not sanctioned
- Suspicious behavior, such as mass audit clearing
- Known malicious threats, such as worm propagation or a password-cracking tool

There are two methods to design a rule in ESA:

- **Rule Builder** is an easy-to-use interface. You provide a meta key and value, then select choices from lists to complete the criteria.
- **Advanced EPL** allows you to write queries in the Event Processing Language. You must know EPL syntax.

If you know EPL, you can use either method. If you do not know EPL, you should use Rule Builder. These topics explain the Rule Builder.


Step 1. Name and Describe the Rule



This topic provides instructions to identify a rule, indicate if it is a trial rule and assign a severity level. When you add a new rule, the first information to provide is a unique name and description of what the rule detects. After you save the rule, this information is displayed in the Rule Library.


Prerequisites

You must have permission to manage rules. See [ESA Permissions](#).


Name and Describe a Rule

1. Go to  (Configure) > **ESA Rules** > **Rules** tab.

- In the **Rule Library**, select   > **Rule Builder**.
The **New Rule** tab is displayed.



The screenshot shows the RSA Rule Builder interface. At the top, there's a navigation bar with 'RSA' logo and various menu items like 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', 'Reports'. Below that, there's a sub-navigation bar with 'LIVE CONTENT', 'SUBSCRIPTIONS', 'INCIDENT RULES', 'INCIDENT NOTIFICATIONS', 'ESA RULES' (highlighted), 'CUSTOM FEEDS', 'EVENT RULES', and 'LOG PARSER RULES'. The main content area is titled 'Rule Builder' and has a subtitle 'Build a rule using drag-and-drop and auto-complete tools.' It contains several sections: 'Rule Name *' (text input), 'Description' (text area), 'Trial Rule' (checkbox checked), 'Memory Threshold' (100 MB), 'Alert' (checkbox checked), 'Severity *' (Low), 'Conditions *' (table with columns: Statement, Occurs, Connector, Correlation Type, Meta, Meta), 'Notifications' (table with columns: Output, Notification, Notification Server, Template), and 'Enrichments' (table with columns: Output, Enrichment Source, ESA Event Stream Meta, Enrichment Source Column Name). There are also 'Group By' and 'Occurs Within' options. The bottom of the interface shows 'RSA NETWITNESS PLATFORM' and '11.5.0.0'.

- Type a unique, descriptive name in the **Rule Name** field.
This name will appear in the Rule Library so be specific enough to distinguish the rule from others.
- In the **Description** field, explain which events the rule detects.
The beginning of this description will appear in the Rule Library.
- By default, new rules are configured as a Trial Rule. A trial rule automatically disables the rule if all trial rules collectively exceed the memory threshold. If you are editing an existing rule, you can select **Trial Rule** to safely test the rule edits.
Use trial rule mode as a safeguard to see if a rule runs efficiently and to prevent downtime caused by running out of memory. For more information, see [Working with Trial Rules](#).
- (This option applies to 11.5 and later.) Enter a **Memory Threshold** for a rule that uses memory, such as a rule that contains windows or pattern matching. If the configured memory threshold is exceeded, the rule gets disabled individually and an error is displayed for that rule on the  (Configure) > ESA Rules > Services tab. The Memory Threshold option works for trial rules and non-trial rules. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.
- (This option applies to 11.3 and later.) Select **Alert** to send an alert to Respond. Clear the checkbox if you do not want to send an alert to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.
- For **Severity**, classify the rule as Low, Medium, High or Critical.

Step 2. Build a Rule Statement

This topic provides instructions to define rule criteria in Rule Builder by adding statements. A statement is a logical grouping of rule criteria in the Rule Builder. You add statements to define what a rule detects.

Example

The following graphic shows an example of a Rule Builder statement.

Every statement contains a key and value. Then, you build logic around the pair by selecting an option in each other field.

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.


Name *

if all conditions are met + -

Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/> event.medium	is	32	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/> event.device_class	is	IDS, Firewall, IPS, Intrusion, Vuln...	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Prerequisites

To build a rule statement, you must know the meta key and the meta value.

For a complete list of meta keys, go to  (Configure) > ESA Rules > Settings > Meta Key References.

Meta entities are not currently supported, such as:

```





fullname.all
eth.all
ip.all
ipv6.all
port.src.all
port.dst.all
dir.path.all
org.all
geoip.all
port.all

```

domain.all
email.all
filename.all
directory.all
checksum.all
param.all
context.all
attack.all
analysis.all
compromise.all
inv.all
outcome.all
ec.all
user.all
host.all
client.all

Caution: If you add meta entities to your rule, they cannot get data from the data sources, so they do not trigger alerts.

Build a Rule Statement

1. Go to  **(Configure) > ESA Rules**.
The Rules tab is displayed by default.
2. In the **Rule Library**, click   **> Rule Builder** or edit an existing Rule Builder rule.
The Rule Builder view is displayed.
3. In the **Conditions** section, click  .
The Build Statement dialog is displayed.

Build a Statement



Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.


Name * Failed login

if all conditions are met

Key	Operator	Value	Ignore Case?	Array?
<input checked="" type="checkbox"/> event.ec_outcome	is	Failure	<input type="checkbox"/>	<input type="checkbox"/>


Cancel Save

4. **Name** the statement. Be clear and specific. The statement name will appear in the Rule Builder.
5. From the drop-down list, select which circumstances the rule requires:
 - if **all conditions** are met
 - if **one of these conditions** are met
6. Specify the criteria for the statement:
 - a. For **Key**, type the name of the **Meta Key**.
 - b. For **Operator** specify the relationship between the meta key and the value you will provide for it. The operator that you use depends on the metadata type. The choices are: is, is not, is not null, is greater than (>), is greater than or equal to (>=), is less than (<), is less than or equal to (<=), is one of (For array type meta), is not one of (For array type meta), contains, not contains, begins with, ends with
 - c. Type the **Value** for the meta key. Do not add quotes around a value. Separate multiple values with a comma.
 - d. The **Ignore Case?** field is designed for use with string and string array values. By choosing the **Ignore Case** field, the query will treat all string text as a lowercase value. This ensures that a rule that searches for the user named Johnson would trigger if the event contains "johnson," "JOHNSON," or "JoHnSoN." (For best rule performance, only use the **Ignore Case** option when necessary.)
 - e. The **Array?** field indicates if the contents of the Value field represent one or more than one value. Select the Array checkbox if you entered multiple, comma-separated values in the **Value** field. For example, "ec_activity is Logon, Logoff" requires you to select the Array checkbox.
7. To use another meta key in the statement, click , select **Add Meta Condition** and repeat step 6.
8. To add a whitelist, click  and select **Add Whitelist Condition**.

9. To add a blacklist, click  and select **Add a Blacklist Condition**.
10. To save the statement, click **Save**.

To Add a Whitelist


You use a whitelist to ensure that specified entities are excluded from triggering the rule. Whitelists can be based on geographic location, in-memory enrichment, or Context Hub list sources. For example, if you want to create a rule that only triggers for IP addresses outside of the US, you can create a whitelist of US IP addresses.

1. After you add a meta condition, click  and select **Add Whitelist Condition**.
2. In the **Enter Whitelist Name** field, select an enrichment source. Any in-memory enrichment, Context Hub list, or a named window in Esper can be used as the source for a whitelist.
3. For the subcondition:
 - a. If you used a GeoIP source for the whitelist, ipv4 is automatically entered for the subcondition. Enter the meta value for the corresponding value field. For example, enter *ipv4 is ip_src* to ensure the GeoIP records are selected based upon the ip_src being found in the GeoIP lookup database. In addition, if you used a GeoIP source for the whitelist, you might want to add a subcondition to specify the geographic region to exclude from the rule results. For example, to specify that the country code must be USA, enter "*CountryCode is US*".
 - b. If you used a Context Hub list for the whitelist, select a column name from the list, then select an operator and enter the meta value for the corresponding value field.

Note: An *Entity* is a specified piece of meta, such as IP address, MAC address, user, host, domain, file name, or file hash.

To Add a Blacklist

You use a blacklist to ensure that specified entities trigger the rule. Blacklists can be based on geographic location, in-memory enrichment, or Context Hub list sources. For example, you can specify that the rule only includes results from Germany.

1. After you add a meta condition, click  and select **Add Blacklist Condition**.
2. In the **Enter Blacklist Name** field, select an enrichment source. Any in-memory enrichment, Context Hub list, or a named window in Esper can be used as the source for a blacklist.
3. For the subcondition:
 - a. If you used a GeoIP source for the blacklist, ipv4 is automatically entered for the subcondition. Enter the meta value for the corresponding value field. For example, enter *ipv4 is ip_src* to ensure the GeoIP records are selected based upon the ip_src being found in the GeoIP lookup database. In addition, if you used a GeoIP source for the blacklist, you might want to add a subcondition to specify the geographic region to include in the rule results. For example, to specify that the rule only includes results for Germany, enter "*CountryCode is DE*".

- b. If you used a Context Hub list for the blacklist, select a column name from the list, then select an operator and enter the meta value for the corresponding value field.

Example: Blacklist

The following statement shows a blacklist statement for a rule that monitors for non-SMTP traffic on TCP destination port 25 containing an executable from countries that are outside of the United States.

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met + -

<input type="checkbox"/>	Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/>	event.service	is not	25	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.tcp_dstport	is	25	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.extension	is	exe,com,vb,vbs,vbe,cmd,bat,ws,ws...	<input type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="checkbox"/>	blacklist.GeoIpLookup				
<input type="checkbox"/>	ipv4	is	event.ip_src	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	countryCode	is not	US	<input type="checkbox"/>	<input type="checkbox"/>

Blacklist conditions can be added to include only those items defined in an enrichment list. Map a list column to an event meta key to join the list to the incoming data stream.

Cancel Save

Statement	Description
service is not 25	The traffic is not SMTP traffic.
tcp_dstport is 25	The traffic is running on TCP port 25.
extension is exe, com,vb,vbs,vbe,cmd,bat,ws,wsf,src,sh	The file extension is an executable.
GeoIpLookup	The blacklist is based on a GeoIPLookup source.
ipv4 is ip_src	The GeoIP records are selected based upon the ip_src being found in the GeoIP lookup database.
countryCode is not US	When looking up the IP address Event.ip_src in the GeoIP database, the record it returns does not contain "US" in the countryCode field.

Example: Strict Pattern Matching and Using the *Is Not Null* Operator

The following example uses the ability to exclude null values and create a strict pattern match to ensure that it returns the expected rule results. The following conditions make up the rule:

Statement	Occurs	Connector	Correlation Type	Meta	Meta
<input type="checkbox"/> Failures	5	followed by			
<input checked="" type="checkbox"/> Success	1	AND			
<input type="checkbox"/> ModifyPassword	1				

Group By: user_dst, ip_src

Occurs Within: 5 minutes Event Sequence: Strict Loose

Rule Condition	Description
Failures	This condition searches for five failed logins with a "followed by" connector, meaning that the condition (Failures) must be followed by the next condition (Success).
Success	This condition searches for one successful login.
ModifyPassword	This condition searches for an instance where the password is modified.
GroupBy: user_dst, ip_src	The GroupBy field ensures that all the previous conditions are grouped by the user_dst meta (the user destination account) and ip_src. This is important to the construction of the rule because the rule attempts to find a case where a user has attempted to log into the same destination account multiple times, finally logged in successfully, and then changed the password. Grouping by ip_src ensures that the user logged in from the same machine attempted to log into an account multiple times. The rule may give unexpected results if you do not group the results.
Occurs within 5 minutes	The time window for the events to occur is five minutes. If the events occur outside of this time window, the rule does not trigger.
Event Sequence: Strict	The event sequence is configured for a strict pattern match. This means that the pattern must match exactly as it is specified with no intervening events. Strict pattern matching allows you to ensure that the Esper engine only generates alerts for rules that exactly match the pattern you want to find. For example, a common rule might be to search for five failed logins followed by a successful login. If you select a loose pattern match, this rule will trigger if there are any number of successful logins between the failed logins. Since the point of the rule is to find frequent <i>and</i> sequential login attempts, a strict match is required to ensure that you get the results you expect.

Note: Each of these conditions is explained in further detail in the sections below.

For each condition, a statement is built in the Rule Builder. The following statement makes up the Failures condition:

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name * Failures

if all conditions are met

<input type="checkbox"/>	Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/>	event.ec_activity	is	Logon	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.ec_outcome	is	Failure	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.user_dst	is not null		<input type="checkbox"/>	<input type="checkbox"/>

Cancel Save

Rule Statement	Description
ec_activity is Logon	Identifies activity that attempts to log on to a system. The Ignore Case field is designed for use with string and string array values. By choosing the Ignore Case field, the query will treat all string text as a lowercase value. You may want to use this field if you are unsure what case may be used when logging a particular event. For best rule performance, only use the Ignore Case option when necessary.
ec_outcome is Failure	Identifies activity outcome logged as "failure."
user_dst is not null	Ensures that the condition is only true if user_dst is populated. The is not null operator allows you to ensure that a field returns a value. You may want to use this field when a rule depends on a particular field returning a value. For example, you want to create a rule that identifies the same user attempting to log into the same destination account multiple times (potentially a password-guessing attack). If the field that represents the user destination account is empty, you don't want the rule to trigger. To ensure the field contains a value, you use the is not null operator.

The following statement makes up the Success condition:

Build a Statement ? X

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met + ⌵ -

	Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/>	event.ec_activity	is	Logon	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.ec_outcome	is	Success	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.user_dst	is not null		<input type="checkbox"/>	<input type="checkbox"/>

Cancel Save

Rule Statement	Description
ec_activity is Logon	Identifies logon activity.
ec_outcome is Success	Identifies a logon that is successful.
user_dst is not null	Ensures that user destination account field must be populated for the condition to be true.

The following statement makes up the ModifyPassword condition:

Build a Statement ? X

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met + ⌵ -

	Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/>	event.user_dst	is not null		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.ec_subject	is	Password	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.ec_activity	is	Modify	<input type="checkbox"/>	<input type="checkbox"/>

Cancel Save

Rule Statement	Description
user_dst is not null	Ensures the user destination account field must be populated for the condition to be true.
ec_subject is Password	Identifies a subject of Password.
ec_activity is Modify	Identifies activity where the password was modified.

Example Results

When the alert fires for the above example rule, you can see that the rule triggered for seven events, and that each event contains a user. You can also see that the events follow a strict pattern: five failed login events, followed by a successful login event, followed by a modification to the account.

The following figure shows the alert in the Respond Alerts List view.

TIME RANGE	CREATED	SEVERITY	NAME	SOURCE	# EVENTS	HOST SUMMARY	INCIDENT ID
Last 5 Minutes	08/25/2017 03:50:43 pm	90	5 Failed Logins Followed By Successful Login Strict...	Event Stream Analysis	7	10.100.33.1 to 7 hosts	

The next figure shows the events in the alert in the Respond Alert Details view.

TIME	TYPE	SOURCE IP	SOURCE PORT	SOURCE HOST	SOURCE MAC	SOURCE USER	DESTINATION IP	DESTINATION P...	DESTINATION HOST	DESTINATION MAC	DESTINATION U
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.33.1				Auser1
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.33.2				Auser1
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.33.3				Auser1
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.33.4				Auser1
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.33.5				Auser1
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.33.6				Auser1
08/25/2017 03:50:40:000 ...	Log	10.100.33.1					10.100.36.78				Auser1

Drilling down into the Investigation view by clicking on the source for one of the events, you can see the case for each of the string values.

The screenshot displays the Malware Analysis interface with the following components:

- Navigation:** "Navigate", "Events", "Malware Analysis".
- Query Bar:** "NWNNodeXLCL51612 - Log Decoder", "Last 5 Minutes", "Query", "Profile", "Detail View", "Actions", "Incidents".
- Filters:** "device.ip exists", "device.disc exists", "device.disc = 85", "device.disc = 85", "Cancel".
- Table:**

Event Time	Event Type	Event Theme	Size	Details
2017-08-25T15:46:11	Log	User.Activity.Failed Logins	137 bytes	<pre> header.id : 0001 level : 6 netname : private src netname : private dst direction : lateral user_dst : Auser3 ec.subject : User ec.activity : Logon ec.theme : Authentication ec.outcome : Failure reference.id : 605004 event.desc : Login denied result : Login denied msg.id : 605004 event.cat.name : User.Activity.Failed Logins device.disc : 85 </pre>
- Buttons:** "Show Additional Meta", "Event Analysis", "Hide Additional Meta", "Event Analysis".

Example: Grouping the Rule Results

The **Group By** field allows you to group and filter rule results. For example, suppose that there are three user accounts; Joe, Jane, and John and you use the **Group By** meta, `user_dst`. The result will show events grouped under the accounts for Joe, Jane, and John.

You can also group by multiple keys, which can further filter rule results. For example, you might want to group by user destination account and machine to see if a user logged into the same destination account from the same machine attempts to log into an account multiple times. To do this, you might group by `user_dst` and `ip_src`.

The following example shows a rule grouped by `user_dst` and `ip_src`.

Rule Builder
Build a rule using drag-and-drop and auto-complete tools.

Rule Name * 5F1S with MultipleGroup by

Description 5 Failed Logins Followed By Successful Login Strict
Group by: Destination User Account and Source IP Address

Trial Rule

Memory Threshold 100 MB

Alert

Severity * Low

Conditions * Investigation

Statement	Occurs	Connector	Correlation Type	Meta	Meta
<input type="checkbox"/> Failed Logins	5	followed by			
<input checked="" type="checkbox"/> Successful Login	1				

Group By user_dst ip_src

Occurs Within 5 minutes Event Sequence Strict Loose

Rule Condition	Description
Failed Logins	Identifies five failed login attempts (must be followed by the next condition; that is, the five failed logins must be followed by a successful login).
Successful Login	Identifies one successful login.
Group By: user_dst and ip_src	Groups the rule results by user_dst (user destination account) and ip_src (IP address of the machine that the user is logging in from). This allows the rule to look for a user logged in from the same machine to the same destination account, resulting in a much more targeted rule result.
Occurs within 5 minutes with a strict pattern match	The events must occur within five minutes, and the pattern matching is strict, meaning it must follow the pattern exactly for the rule to trigger.

Example: Working with Numeric Operators

Numeric operators allow you to write rules against numeric values, such as specifying that a value is greater than, less than, or equal to a specific value. This is useful particularly for cases where you might want to specify a numeric threshold, that is, *payload is greater than 7000*.

The following example attempts to identify a data transfer to a particular destination through the common ports where the transfer size is high and the payload is in a suspicious range.

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met + -

<input type="checkbox"/>	Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/>	event.ip_dst	is	10.10.10.1	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.ip_dstport	is less than or equal	1024	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.size	is greater than or equal	10000	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.payload	is greater than	7000	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.payload	is less than	8000	<input type="checkbox"/>	<input type="checkbox"/>

Rule Statement	Description
ip_dst is 10.10.10.1	The destination port is 10.10.10.1.
ip_dstport is greater than or equal to 1024	The destination port is in a commonly used port range, 1024 or greater.
size is greater than or equal to 10000	The size of the transfer is 10000 or greater, which is a suspiciously large transfer.
payload is greater than 7000	The payload is between 7000 and 8000, which is a suspiciously large payload.
payload is less than 8000	The payload is between 7000 and 8000, which is a suspiciously large payload.

Step 3. Add Conditions to a Rule Statement

This topic provides instructions to add conditions, such as specifying a certain time frame, to a rule statement. When you build a statement, you specify what a rule detects. You add conditions to make further stipulations, such as how many times or when the criteria must occur.

Example

The following graphic shows an example of the conditions for Rule Builder statements. Combined, the statements and conditions comprise the rule criteria.

The screenshot shows the 'Conditions' section of a rule builder interface. It features a table with columns for 'Statement', 'Occurs', 'Connector', 'Correlation Type', 'Meta', and 'Meta'. The 'Success' statement is selected with a checkmark, and its 'Occurs' value is set to 1. The 'Failures' statement has an 'Occurs' value of 5 and is connected to 'Success' with the connector 'followed by'. The 'ModifyPassword' statement has an 'Occurs' value of 1. Below the table, the 'Group By' field is set to 'user_dst' and 'ip_src'. The 'Occurs Within' field is set to 5 minutes, and the 'Event Sequence' is set to 'Strict'.

Statement	Occurs	Connector	Correlation Type	Meta	Meta
<input type="checkbox"/> Failures	5	followed by			
<input checked="" type="checkbox"/> Success	1	AND			
<input type="checkbox"/> ModifyPassword	1				

Group By: user_dst, ip_src


Occurs Within: 5 minutes

Event Sequence: Strict Loose

This rule detects 5 failed logon attempts followed by one successful logon, which could be the sign that someone has hacked into user account. This is the criteria for the rule:

- 5 failed logons are required.
- 1 successful logon must follow the failures
- A password was changed.
- All events must occur within 5 minutes.
- Group alerts by user (user_dst), because steps A and B must be performed on the same user destination account. Also, group by machine (ip_src) to ensure that the user logged in from the same machine attempts to log into an account multiple times.
- The match is a strict pattern, meaning that the pattern must match exactly with no intervening events.

Add Conditions to a Rule Statement

- In the **Conditions** section, select a statement and click .
- For **Occurs**, enter a value to specify how many occurrences are required to meet the rule criteria.
- If you have multiple statements, in the **Connector** field select a logical operator to join one statement to another:
 - followed by
 - not followed by

- AND
 - OR
4. **Correlation Type** applies only to **followed by** and **not followed by**. If you choose a correlation type of SAME, select one meta to correlate on, and if you choose a correlation type of JOIN, select two meta to correlate on. You may want to use JOIN if you are trying to correlate on meta from two different data sources. For example, say you want to correlate an AV alert with an IDS alert. See the examples below for a use case where two meta from different sources are joined.
 5. If events must happen within a specific timeframe, enter a number of minutes in the **Occurs Within** field.
 6. Choose whether the pattern must follow a **Strict** match or a **Loose** match. If you specify a strict match, this means that the pattern must occur in the exact sequence you specified with no additional events occurring in between.

For example, if the sequence specifies five failed logins (F) followed by a successful login (S), this pattern will only match if the user executes the following sequence: F,F,F,F,F,S. If you specify a loose match, this means that other events may occur within the sequence, but the rule will still trigger if all of the specified events also occur. For example, five failed login attempts (F), followed by any number of intervening successful login attempts (S), followed by a successful login attempt might create the following pattern: F,S,F,S,F,S,F,S,F,S which would trigger the rule despite the intervening successful logins.

7. Choose the fields to group by from the dropdown list. The **Group by** field allows you to group and evaluate the incoming events.
 For example, in the rule that detects 5 failed logon attempts followed by 1 successful attempt, the user must be the same, so user_dst is the **Group By** meta key. You can also group by multiple keys. Using the previous example, you might want to group by user and machine to ensure that the same user logged in from the same machine attempts to log in to an account multiple times. To do this, you might group by user_dst and ip_src.

Example

The following graphic shows an example of the conditions for a rule that allow you to evaluate the same entities across multiple devices so you can accomplish complex use cases. For example, you can create a rule that triggers if an IDS (Intrusion Detection System) alert is followed by an AV(Anti-virus) alert for the same workstation. The work station key is not the same between the two (IDS & AV) sources, so you can perform a JOIN in order to evaluate the different entities.

In the IDS alert, the workstation is identified by the source IP address from the IDS alert, and would be compared to the destination IP address from the AV alert.

Statement	Occurs	Connector	Correlation Type	Meta	Meta
<input type="checkbox"/> IDS Check	1	followed by	JOIN	ip_src	ip_dst
<input type="checkbox"/> Antivirus Check	1				

Group By:


Occurs Within: 10 minutes


This is the criteria for the rule:

- A. An IDS alert occurs.
- B. The destination IP address from the AV alert and source IP address for the workstation from the IDS alert are joined to allow you to view the same entities across different sources.
- C. An Antivirus alert follows the IDS alert.

Validate an ESA Rule

You can confirm that an ESA rule generates the expected alerts by testing the rule logic using JSON input data. You can view the alerts in the output, but this test does not send any alert notifications.



1. If you are not already in the rule, go to  **(Configure)** > **ESA Rules** > **Rules** tab and in the **Rule Library**, open the ESA rule that you want to test.
2. Scroll down to the **Test Rule** section.





3. In the **ESA Service** field, select the ESA Correlation service to process the rule. Use the same ESA Correlation service that you plan to use in the ESA rule deployment that contains the rule.
4. In the **Input Data** field, enter the input events to test the rule. Download the events from the Investigate view in JSON format, copy the events, and paste them in this field. You can do this from the Investigate > Navigate view or the Investigate > Events view.

To download the events from the Investigate > Navigate view:

- a. In the main menu, go to **Investigate** > **Navigate** in a new tab, select a data source, and click **Navigate**.
- b. In the Navigate view, click **Load Values** and click a meta value to filter the events.

- c. Save the events as meta in the JSON file format [Save Events > Meta > (name the file) > Export Meta Format: choose JSON].
- d. In the toolbar click the  (**Jobs**) icon and then click **View Your Jobs**.
- e. In the Jobs panel, download your extracted meta, for example: **investigation-2020-May-19-08-30-20.json**.
- f. Go to back to the  (**Configure**) > **ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.

To download the events from the Investigate > Events view:

- a. In the main menu, go to **Investigate > Events** in a new tab.
 - b. In the Events view, enter a query for the ESA rule test.
 - c. Select the events to use and in the **Download** or **Download All** menu, select **Visible Meta as JSON** or **All Meta as JSON**, depending on the size of your selection.
 - d. In the main menu, go to **Dashboards** and in the toolbar click the  (**Jobs**) icon and then click **View Your Jobs**.
 - e. In the Jobs panel, download your extracted meta, for example: **Concentrator_ALL_EVENTS_ALL_META.json**.
 - f. Go to back to the  (**Configure**) > **ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.
5. Click **Test Rule**. The **Output** field shows the output of your rule and you can determine if the results meet your requirements.

Test Rule Get immediate feedback on how the rule runs ^

ESA Service: ESA- ESA Correlation

Input Data:

```

],
"ip.dst": [
  "10.100.10.1"
],
"direction": [
  "inbound"
],
"service.name": [
  "telnet"
],
],

```

Test Rule

Output: Test complete

- ✔ Rule successfully validated
- ✔ Provided input is valid
- ✔ Test ran successfully

Engine Stats

Engine Version	Events Offered	Offered Rate	Runtime Errors
8.4.0	5870	0	-

Rule Stats

Deployed	Statement Fired	Alerts Fired	Events in Memory	Memory Usage	CPU %	Events Matched	Alerted Events	Runtime Errors	Debug Logs
✔	0	5870	0	0	100	5870	Details...	-	Details...

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.

Field	Description
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Working with Rules



This topic discusses additional procedures you can perform on rules. You may want to perform any of the following procedures:

- [Edit, Duplicate or Delete a Rule](#)
- [Filter or Search for Rules](#)
- [Import or Export Rules](#)


Edit, Duplicate or Delete a Rule


This topic provides instructions to edit, duplicate, or delete an Event Stream Analysis (ESA) rule. When you edit a rule, ESA applies the updated criteria going forward. No changes are made to previously generated alerts.

Edit a Rule

1. Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.
2. In the **Rule Library**, select the rule you want to edit and click .
Depending on the rule type, the respective rule tab is displayed.
3. Modify the required parameters.
4. Click **Save**.


Duplicate a Rule


1. In the **Rule Library**, select the rule you want to duplicate and click .
2. The Duplicate a Rule dialog is displayed. The system adds **Copy of** in front of the rule name.




3. In the **Name** field, type a unique name for the duplicate rule and click **OK**.
A duplicate rule with the new name is added to the Rule Library.

Delete a Rule

1. Go to  (**Configure**) > **ESA Rules** > **Rules**.
The Rules tab is displayed.



- In the Rule Library, select one or more rules and click  .
A warning dialog is displayed.
- Click **Yes**.
A confirmation message that the rule is deleted successfully is displayed and the selected rule is deleted from the Rule Library.



Filter or Search for Rules

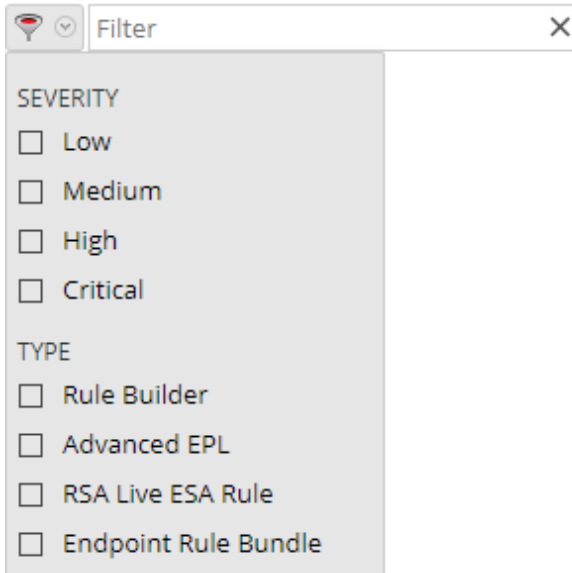
This topic shows analysts how to specify the type of rules that display in the Rule Library.

Prerequisites

Make sure that you understand the Rule Library view components. For more information, see [Rule Library Panel](#).


Filter Rules

- Go to  (**Configure**) > **ESA Rules**.
The Rules tab is displayed by default.
- In the **Rule Library** panel toolbar, click  and select the severity and type of rules that you would like to appear in the Rule Library list. The following figure shows the Filter drop-down list.



The selected rule types appear in the list.

Search for Rules

1. Go to  (**Configure**) > **ESA Rules**.
The Rules tab is displayed by default.
2. In the **Rule Library** panel toolbar, type a rule name in the Filter field.
The Rule Library panel lists the rules that match the names entered in the Filter field.


Import or Export Rules

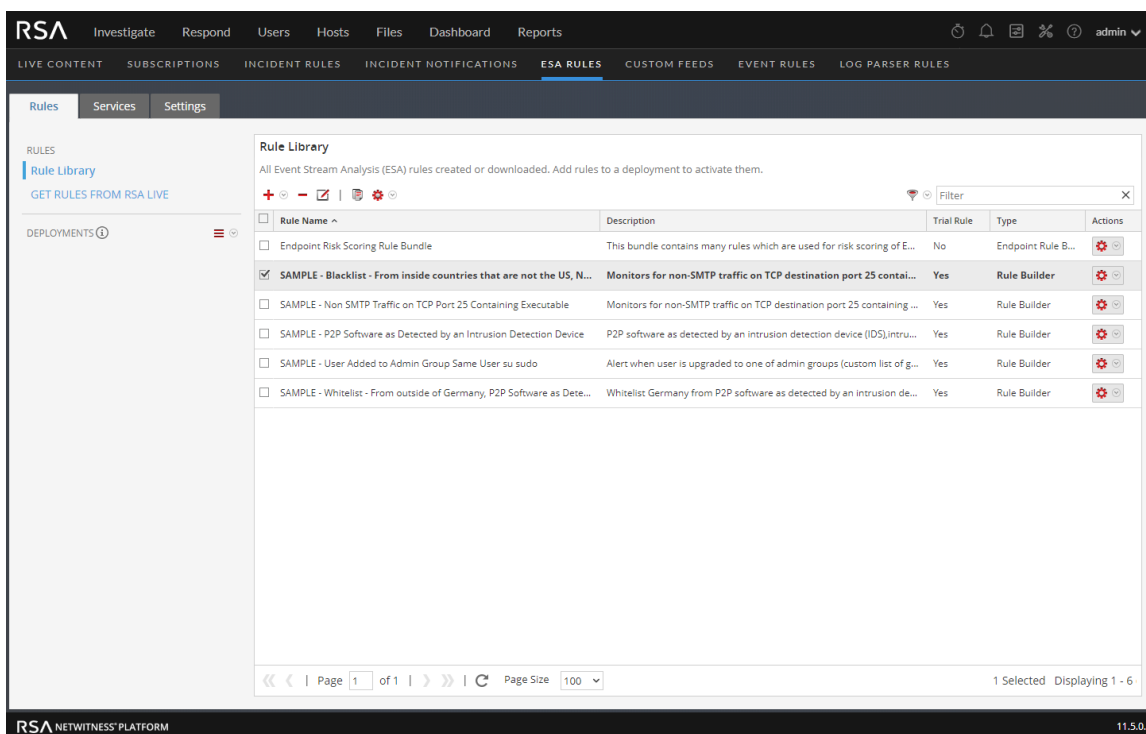
The topic provides instructions to import ESA rules from a NetWitness instance and to export ESA rules to your hard drive so you can keep a local copy.


If you exported a rule in an earlier version of NetWitness, the following conditions apply when you import the rule in version 10.5 or later:


- Exported in version 10.3 – You cannot import rules to version 10.5 or later.
- Exported in version 10.4 – You can import rules to version 10.5 or later.

Import ESA Rules

1. Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.




2. In the **Rules Library** toolbar, select  > **Import**.
The Import ESA Rules dialog is displayed.



3. Click **Browse** to browse and select the file containing the ESA rules.
4. Click **Import**.

Export ESA Rules

1. Select an ESA rule or multiple rules and select  > **Export** in the Rule Library toolbar.
A warning dialog is displayed.
2. Click **Yes**.
The Export Rules dialog is displayed.
3. In the **Enter File Name** field, type a filename for the file with the ESA rules and click **Export**.
The file is exported as a binary file to your machine.

Note: The binary file cannot be edited.

Choose How to be Notified of Alerts

This topic explains the different notification methods and how to add a notification method to a rule. Administrator, SOC Manager or DPO role permissions are required for all tasks in this section.

When a rule triggers an alert, ESA can send a notification in the following ways:

- Email
- Syslog
- Script

To configure a notification, you configure these components:

- **Notification Server:** The notification server is the source of the notifications. After you configure a notification server, you can add it to a rule. When the rule triggers an alert, the rule will use that server to send alert notifications.
- **Notifications:** These are the outputs (destinations) of the notifications, which can be email, script, and Syslog. When you design a rule, you can specify the notification for an alert.
- **Templates:** The message format of an alert notification is defined in a template.

If you use an ESA rule that has an enrichment, such as a Context Hub list, you must create a custom template. You can duplicate a default template and adjust it for your enrichment. For more information, see [Troubleshoot ESA Rules](#). For information on creating a custom template, see "Configure Meta Keys as Arrays in ESA Correlation Rules" in the *System Configuration Guide*.

Go to the [NetWitness Documents - All Versions](#) to find all RSA NetWitness Platform 11.x documents.

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

Alert suppression and alert rate regulation are two features that Event Stream Analysis provides. Alert suppression ensures that multiple emails are not sent out for the same alert. For example, consider a rule to detect failed user logins. If you set the alert suppression to three minutes, you will see only the alerts generated in that time frame. This is fewer than the number of alerts you would see without alert suppression. Some alerts can be duplicates. With alert suppression, emails are not sent for duplicate alerts. This ensures the inbox is not flooded with redundant alert notifications.

Alert rate regulation is a preventive measure to ensure that alerts from misconstrued rules do not flood the system. This ensures that ESA does not send more than the configured limit of emails within one minute.

Notification servers, notifications, and templates are configured in the Administration System view. For more information, see "Configure Notification Servers", "Configure Notification Outputs", and "Configure Templates for Notifications" in the *System Configuration Guide*.

Notification Methods

When a rule triggers an alert, ESA can send a notification in the following ways:

- Email
- Syslog

- Script

Note: ESA SNMP notifications are not supported for NetWitness 11.3 and later.

Email Notifications

ESA Correlation can send notifications to users through email about various system events.

To configure these email notifications, you need to:

- Configure the SMTP email server as an output provider. For instructions, see "Configure the Email Settings as Notification Server" in the *System Configuration Guide*.
- Set up an email account to receive notifications. For instructions, see "Configure Email as a Notification" in the *System Configuration Guide*.
- Configure a template for email notification. For instructions, see "Configure Global Notifications Templates" in the *System Configuration Guide*.

Syslog

Event Stream Analysis can send events and consolidate logs in Syslog format to a Syslog server.

To configure these Syslog notifications, you need to:

- Configure Syslog server settings as an output provider. For instructions, see "Configure a Syslog Notification Server" in the *System Configuration Guide*.
- Configure Syslog message format as an output action. For instructions, see "Configure Syslog as a Notification" in the *System Configuration Guide*.
- Configure a template for Syslog. For instructions, see "Configure Global Notifications Templates" in the *System Configuration Guide*.

Script Alerter

Apart from the alert notifications ESA allows users to run scripts in response to ESA alerts.

Scripts enable you to do custom integration with applications that exist in your environment. For example, if you want to open an incident ticket from an application when a specific alert is triggered, Script Alerter lets you write a script that calls the application API and has ESA invoke it when the specific ESA rule is triggered. You can configure a FreeMarker template to define what details you want to extract from the output of the ESA rule and pass it as command line arguments to the script.

To use the Script Alert, you need to:

- Configure the user identity and other details that are required to execute the script. For instructions, see "Configure Script as a Notification Server" in the *System Configuration Guide*.
- Define the Script. For instructions, see "Configure Script as a Notification" in the *System Configuration Guide*.
- Configure a template for the script. For instructions, see "Configure Global Notifications Templates" in the *System Configuration Guide*.

Add Notification Method to a Rule

This topic tells administrators how to add a notification, such as email, to a rule. ESA uses the notification method when it generates an alert for an event that meets rule criteria.


You add a notification to a rule so ESA can let you know when a rule triggers an alert. Although the notification fields are not required, it is a best practice to add a notification to a rule.

When you add a notification method to a rule, you select the following information:

- Output
- Notification
- Notification Server
- Template




Prerequisites


- Your role must have permission to manage rules.
- The rule must exist.
- The notification method must be configured with a supported server and template:

Go to  (Admin) > System > Global Notifications.

For detailed procedures, see the *System Configuration Guide*.


Add a Notification Method to a Rule

1. Go to  (Configure) > ESA Rules > Rules tab.
2. In the **Rule Library**, click  to add a new rule or select an existing rule and click . Depending on the rule type, the Rule Builder or Advanced EPL tab is displayed. The Notifications section is the same for both tabs.



Output	Notification	Notification Server	Template
No parameters to edit.			

Output Suppression of every minutes

3. Click  and select the **Output** for the alert:
 - Email
 - SNMP (This option is not supported in NetWitness 11.3 and later.)
 - Syslog
 - Script

4. Double-click the **Notification** field and select the name of a previously configured output.
For example, Level 1 Analyst could be the name of an email notification that goes to the L1-Analysts email distribution group.
5. Double-click the **Notification Server** field and select the server that sends the notification.
6. Double-click the **Template** field and select a format for the alert.
The following figure shows the settings for a Syslog notification.

Output	Notification	Notification Server	Template
<input checked="" type="checkbox"/> SYSLOG	Local_SysLog	localhost-514	Default Syslog Template


Output Suppression of every minutes

7. If you want to specify frequency, select **Output Suppression**, then enter the number of **minutes**.
8. If you want to add another notification, repeat steps 3-7.
9. Click **Save**.
When ESA generates an alert for an event that matches the rule criteria, you will be notified of the alert via each notification method added to the rule.

Add a Data Enrichment Source

This topic tells how to add a previously configured enrichment source to a rule. When ESA creates an alert, information from the source gets included in it.

Enrichments provide the ability to include contextual information into correlation logic and alert output. Without enrichments, all information included in an ESA alert is from a Core service. With enrichments, you can request for look ups into a variety of sources and include the results into the outgoing alerts. The following figure illustrates the enrichment feature.



Enrichment configuration is made up of two logical units:

- Enrichment Sources – These are data stores of contextual information.
- Enrichment Connections – These act as connectors between alert meta and source columns.

ESA allows you to make connections between Event Processing Language (EPL) statements and enrichment sources. Once the connections are established, the system joins the selected fields from the alert output with the information in the sources and uses the matching data to enrich the alert that is sent out. ESA can connect with the following sources:

- Esper Named Windows
- MaxMindGeoIP Database

Note: The geoIP enrichment source can neither be created nor deleted. It is provided out of the box to the user.

Example Rule with Enrichments

The following example rule illustrates how ESA enrichments can enhance alerts.

```
@RSAAlert @Name("simple") SELECT * FROM Event(ec_theme='Login Failure')
```

This rule generates an alert for every logon failure and thus if the following (simplified) event stream is received at ESA:

sessionid	ec_theme	username	ip_src	ip_dst	host_dst
1	Login Success	dshrute	23.xx.23x.16		
2	Login Failure	jhalpert	23.xx.23x.16	31.1x.x9.1x8	www.facebook.com

An alert without an enrichment with the following constituent events might be generated in response to the second session:

```
{
  "events": [
    {
      "username": "jhalpert",
      "host_dst": "www.facebook.com",
      "ip_dst": "31.1x.x9.1x8",
      "sessionid": 2,
      "ec_theme": "Login Failure",
      "esa_time": 1406148964130,
      "ip_src": "23.xx.23x.16"
    }
  ]
}
```

The JSON output shows all the information available for inclusion into an ESA notification using an appropriate FreeMarker template. For instance, the template expression `${events[0].username}` would evaluate to `jhalpert`.

With enrichments, the same deployment, with the same event stream, can generate the alert shown below.

```
{"events": [
  {
    "username": "jhalpert",
    "host_dst": "www.facebook.com",
    "GeoIpLookup": [
      {
        "city": "Cambridge",
        "longitude": -71,
        "countryCode": "US",
        "areaCode": 617,
        "metroCode": 506,
        "region": "MA",
        "dmaCode": 506,
        "ipv4Obj": "/23.xx.23x.16",
        "countryName": "United States",
        "postalCode": "02142",
        "ipv4": "23.xx.23x.16",
        "latitude": 42,
        "organization": "Verizon Business"
      }
    ],
    "orgchart": [
      {
        "supervisor": "mscott",
        "name": "James Halpert",
        "extension": 3692,
        "location": "Scranton",

```

```

        "department": "Sales",
        "id": "jhalpert"
    }
  ],
  "ip_dst": "31.1x.x9.1x8",
  "sessionid": 2,
  "LoginRegister": [
    {
      "username": "dshrute",
      "ip_src": "23.xx.23x.16"
    }
  ],
  "ec_theme": "Login Failure",
  "esa_time": 1406155218912,
  "ip_src": "23.xx.23x.16"
}
}]

```

The system pulls contextual data to make the alert more meaningful.

To include the name of the supervisor and the name of the user with the last successful login in the ESA notification, this example includes the following template expressions:

`${events[0]["orgchart"][0].supervisor}` gives the name of the supervisor of the employee in the alert and `${events[0]["LoginRegister"][0].username}` gives the name of the user with the last successful logon from the same `ip_src` (using a stream based Named Window).

Enrichment Sources

This topic explains options for adding an external data source to provide additional information in alerts. Enrichment sources provide additional information in alerts. For example, an in-memory table can provide a full name, title, office location, and employee number if a user matches rule criteria. The following types of enrichment sources are available:

- Context Hub List (Preferred)
- In-Memory Table (Ad hoc only)
- GeoIP

Note: Database, Database Connection, Warehouse Analytics, and Recurring In-Memory Tables as enrichment sources are not supported for the ESA Correlation service in NetWitness 11.3 and later.

It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources. You can share Context Hub List enrichment sources across the NetWitness. You can only use the In-Memory Table with ESA. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources.

Note: The geoIP enrichment source can neither be created nor deleted. It is provided out of the box to the user.

Configure a Context Hub List as an Enrichment Source

Note: The information in this topic applies to NetWitness Version 11.1 and later.

This topic provides instructions on how to configure a Context Hub list as an enrichment source for ESA. Once a Context Hub list is added as an enrichment source, analysts can use the configured list as a statement condition when creating an ESA rule. Any changes made to the list from within Context Hub are automatically reflected in the enrichment source in real-time. For example, you could create a list of IP addresses in Context Hub and then use that list as either a blacklist or whitelist as part of a correlation rule condition. Any subsequent changes made to the IP list in Context Hub will be reflected in the enrichment source in real-time, to ensure the correlation rule operates with a constantly updating set of information.


Prerequisites

Before configuring a Context Hub list as an enrichment source, the list must first be created as a data source in Context Hub. Any list created in Context Hub is supported and the lists may contain string or numeric values, including IP addresses. For information on creating a list as a data source in Context Hub, see the *RSA NetWitness Context Hub Configuration Guide*.

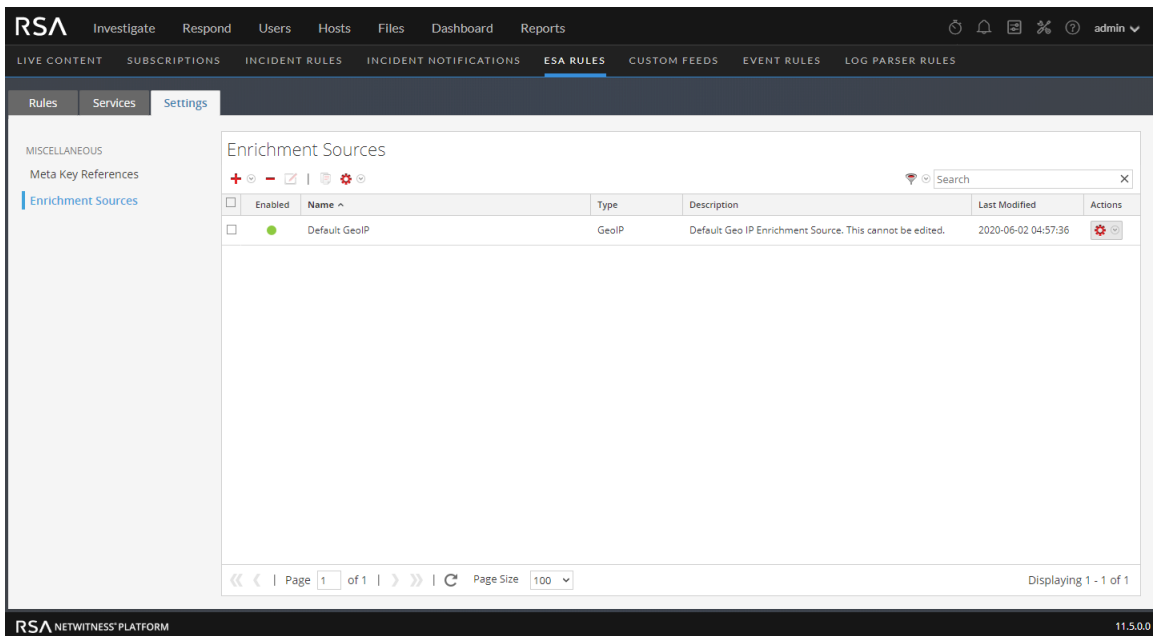
Caution: When creating a Context Hub list for use as an enrichment source, the list name and its field names cannot include any spaces or special characters, or start with a number. If you do not follow this naming convention, when you attempt to add the list as an enrichment source in ESA, an error message will be displayed and you will not be allowed to add the list.

IMPORTANT: If you rename a Context Hub list or recreate the Context Hub list with the same name, update the ESA rules that use that Context Hub list, and then redeploy the ESA rule deployments that contain those rules.

Configure a Context Hub List as an Enrichment Source

1. Go to  (Configure) > ESA Rules > Settings tab.
2. In the options panel, select **Enrichment Sources**.

The Enrichment Sources panel is displayed.



3. From the   drop-down menu, select **Context Hub**.

4. Select **Enable** to enrich alerts with a Context Hub list. This is selected by default. If disabled, the alerts will not be enriched with the configured Context Hub list.
5. Select the desired Context Hub list from the **Select List** drop-down menu of pre-configured lists.
6. (Optional) In the **Description** field, type a brief description about the selected Context Hub list. The text entered here is displayed on the Enrichment Sources panel.

7. In the **Columns** field, all columns included in the selected Context Hub list are listed. Click to enable or disable the columns in the list that you wish to include when using this list as an enrichment source in an alert.
8. (Optional) Click to enable the **Page To Local Store** option. This option is useful if you have a very large list and performance is affected. If this is the case, enabling this option will write a copy of the Context Hub list to the local disk to improve performance.
9. Click **Save**.
The Context Hub list is configured. You can now add it to an ESA rule as part of a condition statement as either a blacklist or a whitelist condition.

The following figure illustrates adding a Context Hub list as part of a condition statement. In this example, a context Hub list named "multicolumnlist" was added as a blacklist condition. The list contains two columns, SourceCity and DestinationCity. The next step would be to select one of the column names as the subcondition and then specify the operator and enter the meta value for the corresponding value field.

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met + -

	Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/>	event.city_src	is not null		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	event.city_dst	is not null		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	blacklist.multicolumnlist				
<input checked="" type="checkbox"/>	<input type="text" value="SourceCity"/>	is	Select...	<input type="checkbox"/>	<input type="checkbox"/>
	DestinationCity				

Blacklist conditions can be added to include only those items defined in an enrichment list. Map a list column to an event meta key to join the list to the incoming data stream.

For complete details for adding a whitelist or blacklist to a condition statement, see [Step 2. Build a Rule Statement](#).

To add a Context Hub list as a condition to an existing rule, select to edit the desired rule in the Rule Library, then add a condition in the Conditions section and select to add a whitelist or blacklist condition to the new condition statement.

Configure an In-Memory Table as an Enrichment Source

This topic provides instructions on how to configure an in-memory table. When you configure an in-memory table, you upload a .CSV file as an input to the table. You can associate this table with a rule as an enrichment source. When the associated rule generates an alert, ESA will enrich the alert with relevant information from the in-memory table.

For example, a rule could be configured to detect when a user tries to download freeware and to identify the person by user ID in the alert. The alert could be enriched with additional information from an in-memory table that contains details such as full name, title, office location and employee number.

Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see [Configure a Context Hub List as an Enrichment Source](#).

Prerequisites

- The column name in the .CSV file cannot have whitespace characters.
For example *Last_Name* is correct, and *Last Name* is incorrect.
- The .CSV file must begin with a header line that defines fields and types.
For example, *address string* would define the header field as *address*, and the type as *string*.

The following shows a valid .CSV file represented as a .CSV and as a table.

	A	B	C
1	address string	criticality integer	department string
2	172.31.110.27	1	SALES
3	172.31.110.28	10	ACCOUNTING
4	172.31.110.29	20	SALES
5			


```

ServerCriticality.csv
address string,criticality integer,department string
172.31.110.27,1,SALES
172.31.110.28,10,ACCOUNTING
172.31.110.29,20,SALES

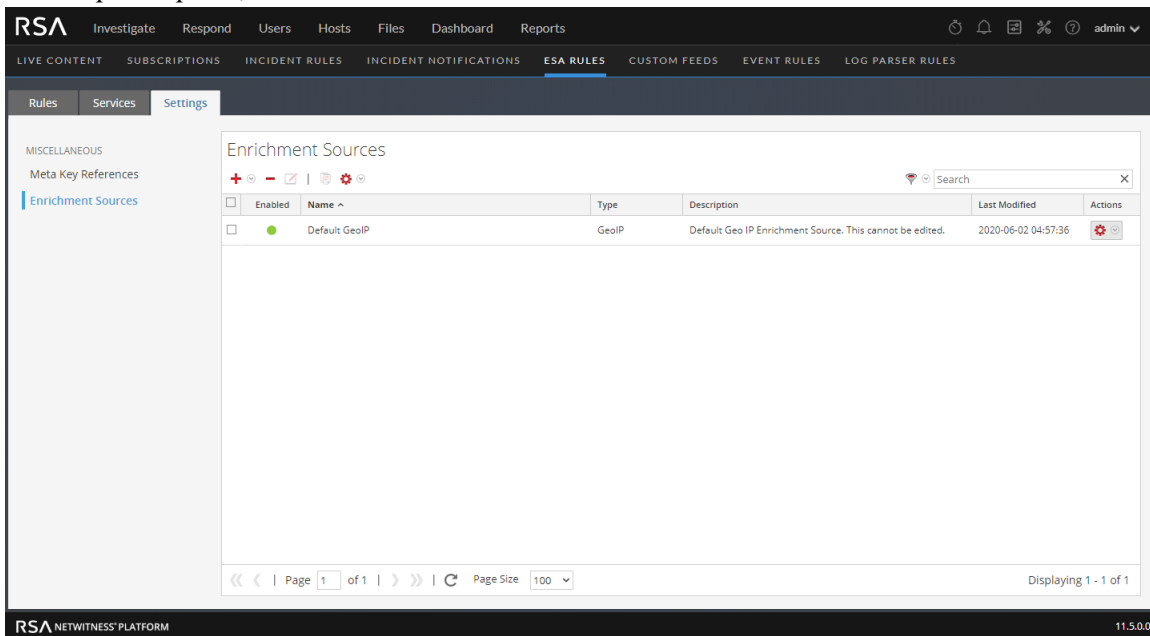
```

Configure an Ad hoc In-Memory Table

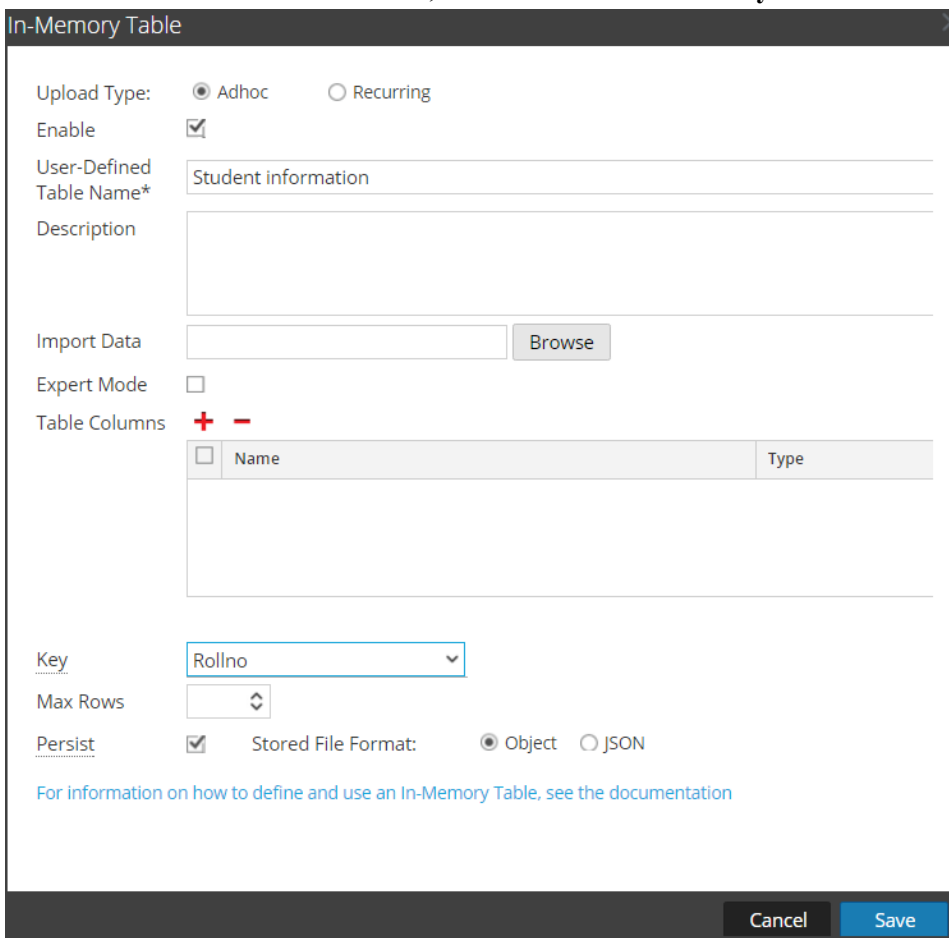
Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see [Configure a Context Hub List as an Enrichment Source](#).

1. Go to  (**Configure**) > **ESA Rules**.
The Configure view is displayed with the ESA Rules tab open.
2. Click the **Settings** tab.

- In the options panel, select **Enrichment Sources**.




- In the **Enrichment Sources** section, click > **In-Memory Table**.



5. Describe the in-memory table:
 - a. Select **Ad hoc**.
 - b. By default, **Enable** is selected. When you add the in-memory table to a rule, alerts will be enriched with data from it.
If you add an in-memory table to a rule but do not want alerts to be enriched, deselect the checkbox.
 - c. In the **User-Defined Table Name** field, type a name, such as Student Information, for the in-memory table configuration.

Note: Do not use any Esper keyword as **User-Defined Table Name** since this causes an error while using this enrichment in the ESA Rule. For Esper keywords, see [Reserved keywords](#).

- d. If you want to explain what the enrichment adds to an alert, type a **Description** such as:
When an alert is grouped by Rollno, this enrichment adds student information, such as name and marks.
6. In the **Import Data** field, select the .CSV file that will feed data to the in-memory table.
7. If you want to write an EPL query to define an advanced in-memory table configuration, select **Expert Mode**.
The Table Columns are replaced by a **Query** field.
8. In the **Table Columns** section, click  to add columns to the in-memory table.
9. If a valid file is selected in the Import Data field, the columns populate automatically.

Note: If you selected Expert mode, a Query field is displayed instead of Table Columns.

10. In the **Key** drop-down menu, select the field to use as the default key to join incoming events with the in-memory table when using a CSV-based in-memory table as an enrichment. By default, the first column is selected. You can also later modify the key when you open the in-memory table in enrichment sources.
11. In **Max Rows** drop-down menu, select the number of maximum number of rows that can reside in the in-memory table at a particular instance.
12. Select **Persist** to preserve the in-memory table on disk when the ESA service stops and to repopulate the table when the service restarts.
13. In **Stored File Format** field, do one of the following:
 - Select **Object**, if you want to store the file in a binary format.
 - Select **JSON**, if you want to store the file in a text format.
By default, **Object** is selected.
14. Click **Save**.
The adhoc in-memory table is configured. You can add it to a rule as an enrichment or part of the rule condition. See [Add an Enrichment to a Rule](#).

When you add an in-memory table, you can add it to a rule as an enrichment or as a part of the rule condition. For example, the following rule uses an in-memory table as a part of the rule condition to create a whitelist, and it also uses an in-memory table of details in the user_dst file to enrich the alert that is displayed.

The rule shows the in-memory table as a whitelist rule condition:

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met

Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/> event.user_dst	is not null		<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> whitelist.User_list				
<input type="checkbox"/> Username	is	event.user_dst	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Whitelist conditions can be added to exclude only those items defined in an enrichment list. Map a list column to an event meta key to join the list to the incoming data stream.

Cancel Save

Next, the alert is enriched with the User_list in-memory table:

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
<input checked="" type="checkbox"/> In-Memory Table	User_list	user_dst	Username

Therefore, the user_dst in-memory table is used to create a whitelist, and it is also used to enrich the data in the alert if the alert is triggered.

Add a Recurring In-Memory Table

Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see [Configure a Context Hub List as an Enrichment Source](#).

It is preferable to use Context Hub List enrichment sources for ESA rules instead of In-Memory Table enrichment sources. You can share Context Hub List enrichment sources across the NetWitness Platform. You can only use the In-Memory Table with ESA.

Note: Database, Database Connection, Warehouse Analytics, and Recurring In-Memory Tables as enrichment sources are not supported for the ESA Correlation service in NetWitness 11.3 and later.




Add an Enrichment to a Rule

This topic tells how to add a previously configured enrichment source to a rule. When ESA creates an alert, information from the source gets included in it.

Adding an enrichment to a rule allows you to request for look ups into a variety of sources and include the results in the outgoing alerts, giving you a more detailed alert. This procedure requires role permissions for Administrator, DPO, and SOC Manager.

Note: This procedure does not apply to adding a Context Hub list as an enrichment to a condition statement in an existing rule. For information see [Configure a Context Hub List as an Enrichment Source](#).

To add an enrichment to a rule:

1. Go to  (Configure) > ESA Rules.
2. In the **Rule Library** view, do one of the following:
 - Double-click a rule.
 - Select a rule and click  in the **Rule Library** toolbar. The Rule Builder panel is displayed in a new NetWitness tab.
3. In the **Enrichments** section, click  and select any of the following enrichment types:
 - In-Memory Table
 - GeoIP

Note: If you use a GeoIP source, ipv4 is automatically populated, and is not editable.

The enrichment types that you have selected are displayed in the table.

4. For the added enrichment type, perform the following:
 - In the **Output** column, select the type that you have configured.
 - In the **Enrichment Source** drop-down list, select the enrichment source defined.
 - In the **ESA Event Stream Meta** field, type the event stream meta key whose value will be used as one operand of join condition.

Enrichments		Settings		
Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name	
<input type="checkbox"/> In-Memory Table	Select Enrichment Source	Enter Meta	Enter Column Name	
<input checked="" type="checkbox"/> GeoIP	Select Enrichment Source	Enter Meta	ipv4	

- In the **Enrichment Source Column Name** field, type the enrichment source column name whose value will be used as another operand of the join condition.
5. Select **Debug**. This adds an @Audit('stream') annotation to the rule. This is useful when debugging the Esper rules.
 6. Click **Show Syntax** to test if the defined ESA rule is valid.
 7. Click **Save**.

For details on parameters and their descriptions, see [Rule Builder Tab](#).

Deploy Rules to Run on ESA

This section explains how an ESA Rule Deployment works and how to set up a deployment to run a group of ESA rules. Administrator, SOC Manager, or Data Privacy Officer role permissions are required for all procedures in this section.

To create an ESA rule deployment, you need to perform the steps described in [ESA Rule Deployment Steps](#).

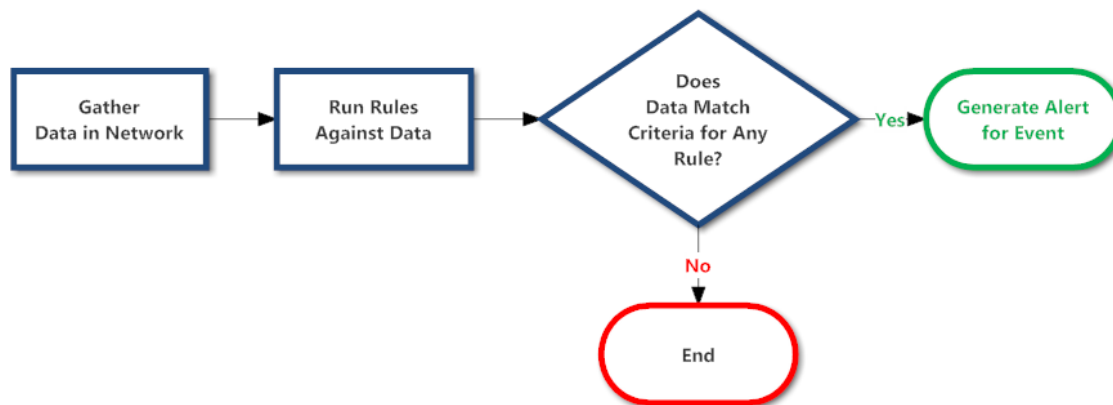
How an ESA Rule Deployment Works

An ESA rule deployment consists of an ESA service, one or more data sources, and a set of ESA rules. When you deploy rules, the ESA service runs them to detect suspicious or undesirable activity in your network. Each ESA rule detects a different event, such as when a user account is created and deleted within one hour.

The ESA service performs the following functions:

1. Gathers **data** in your network
2. Runs ESA **rules** against the data
3. Applies rule **criteria** to data
4. Generates an **alert** for the captured event

The following graphic shows this workflow:



In addition, you may want to perform other steps on your deployment, such as replacing an ESA service, changing a data source, editing or deleting a rule from the deployment, renaming or deleting the deployment, or showing updates to the deployment. For descriptions of these procedures, [Additional ESA Rule Deployment Procedures](#).

ESA Rule Deployment Steps

This topic explains how to add an ESA rule deployment, which includes an ESA service with its associated data sources and a set of ESA rules. You can add an ESA rule deployment to organize and manage ESA services and rules. Think of the deployment as a container for these components:

- An ESA service
- One or more data sources (This is available in version 11.3 and later.)
- A set of ESA rules

For example, if you add a Spam Activity deployment it could include an ESA London service, Concentrators with the appropriate data, and a set of ESA rules to detect suspicious email activity.

Note: An ESA rule deployment can have only one ESA service. You can, however, use the same ESA service in multiple deployments. In NetWitness version 11.2 and earlier, the ESA service is the Event Stream Analysis service. In version 11.3 and later, it is the ESA Correlation service.

To add an ESA rule deployment, you need to complete the following procedures:

- [Step 1. Add an ESA Rule Deployment](#)
- [Step 2. Add an ESA Service](#)
- [Step 3. Add Data Sources](#)
 - [\(Optional\) Add a Data Source Filter](#)
- [Step 4. Add and Deploy Rules](#)



Step 1. Add an ESA Rule Deployment

Prerequisites

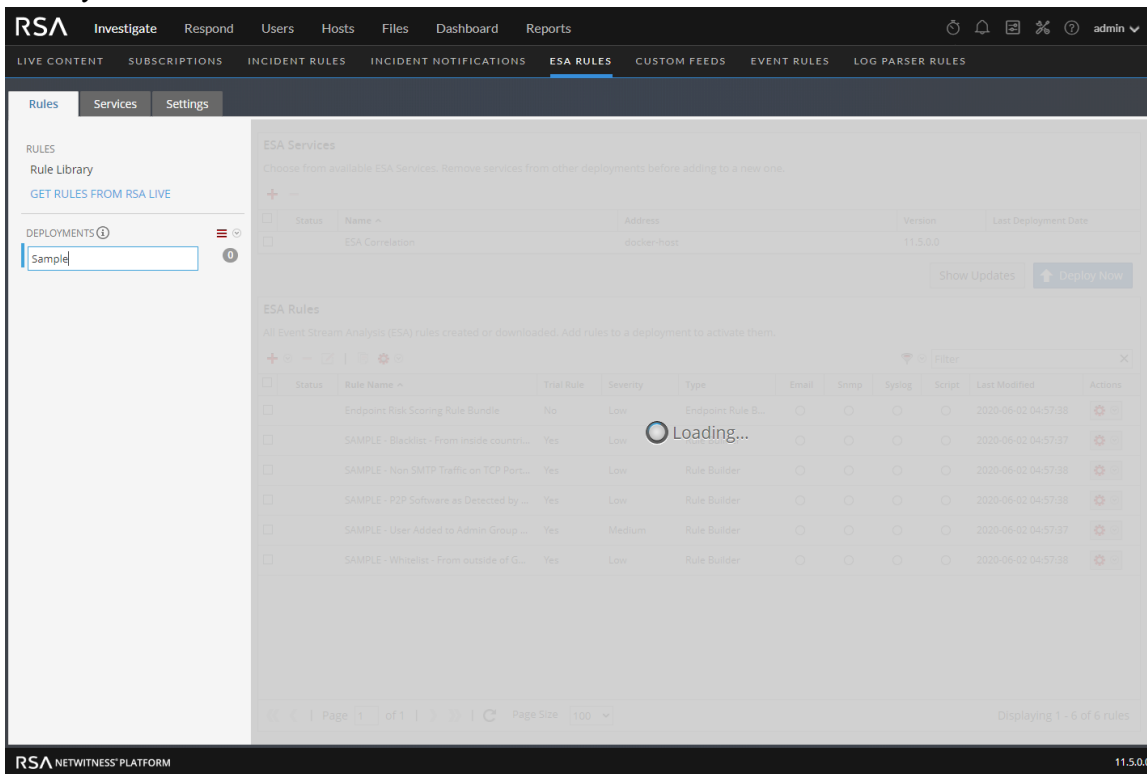
The following are required to add an ESA rule deployment:


- The ESA service must be configured on the host.
- Rules must be in the Rule Library. See [Add Rules to the Rule Library](#).

To add an ESA rule deployment:

1. Go to  **(Configure)** > **ESA Rules**.
The Rules tab is displayed.
2. In the options panel on the left, next to Deployments, select  > **Add** and type a **name** for the deployment. The naming convention is up to you. For example, it could indicate the purpose or

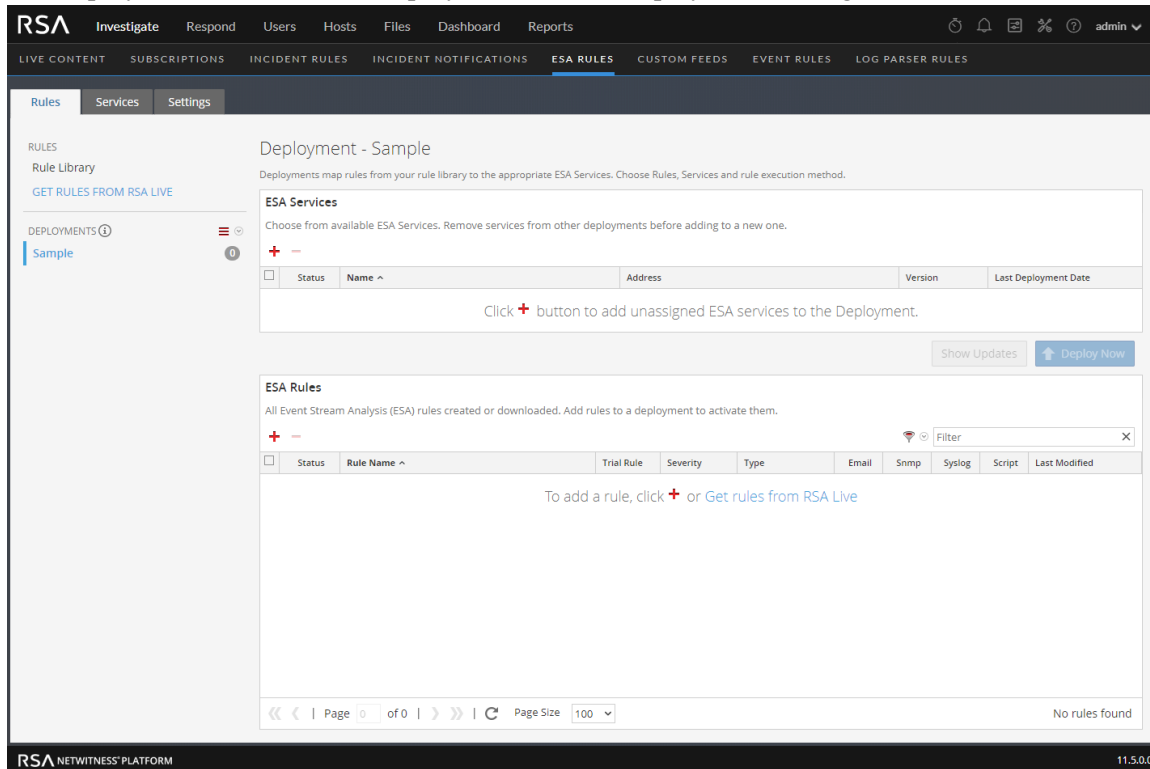
identify an owner.



In NetWitness 11.3 and later, the deployment names that you choose appear on the deployment tabs in the  (Configure) > ESA Rules > Services tab.

3. Press **Enter**.

The deployment is added. The Deployment view is displayed on the right.



Step 2. Add an ESA Service

The ESA service in an ESA rule deployment gathers data in your network and runs ESA rules against the data. The goal is to capture events that match rule criteria, then generate an alert for the captured event.


An ESA rule deployment can have only one ESA service. You can, however, use the same ESA service in multiple deployments. For example, ESA London could be in these deployments simultaneously:

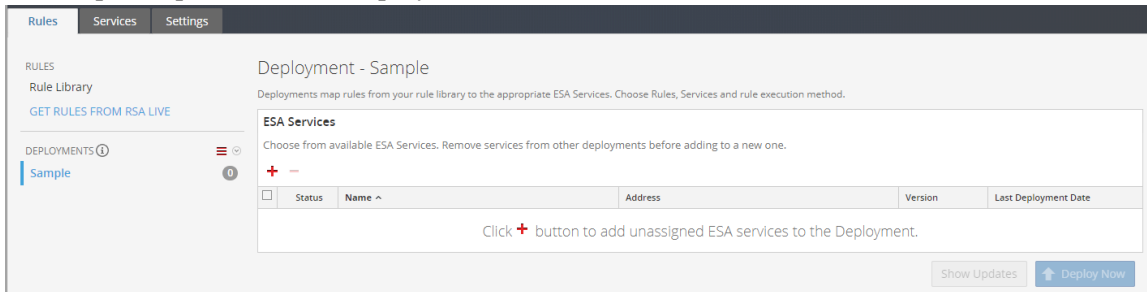
- Deployment EUR, which includes one set of rules
- Deployment CORP, which includes another set of rules.


Changes made to an ESA rule deployment do not take effect until you click **Deploy Now**. For example, Deployment EUR could include the ESA London service and a set of 25 rules. If you replace the ESA London service with the ESA Paris service, the next time you deploy Deployment EUR, the 25 rules will be removed from ESA London and added to ESA Paris.

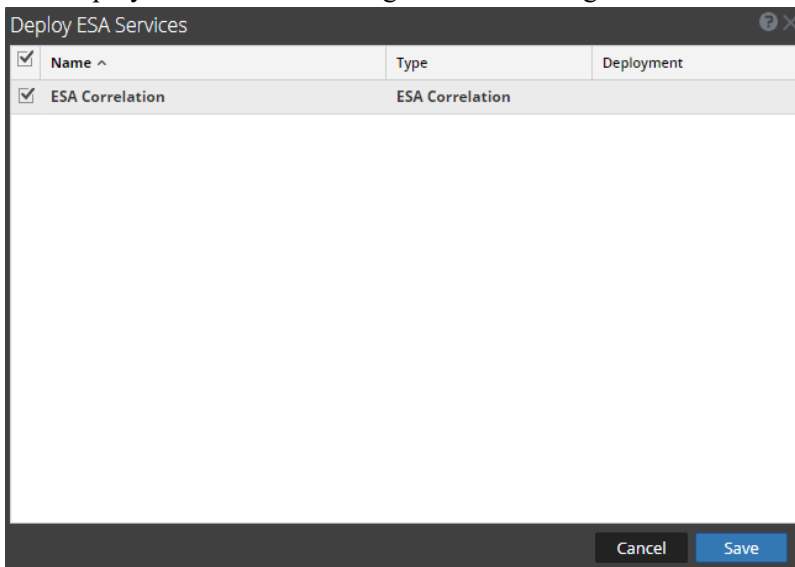
Deleting an ESA rule deployment immediately removes the rules from the ESA service. If an ESA service is not part of any deployment, the ESA service does not have any rules.

To add an ESA service:

1. Go to  (Configure) > **ESA Rules** > **Rules** tab.
2. In the options panel, select a deployment:



3. In the **Deployment** view, click  in **ESA Services**.
The Deploy ESA Services dialog lists each configured ESA.





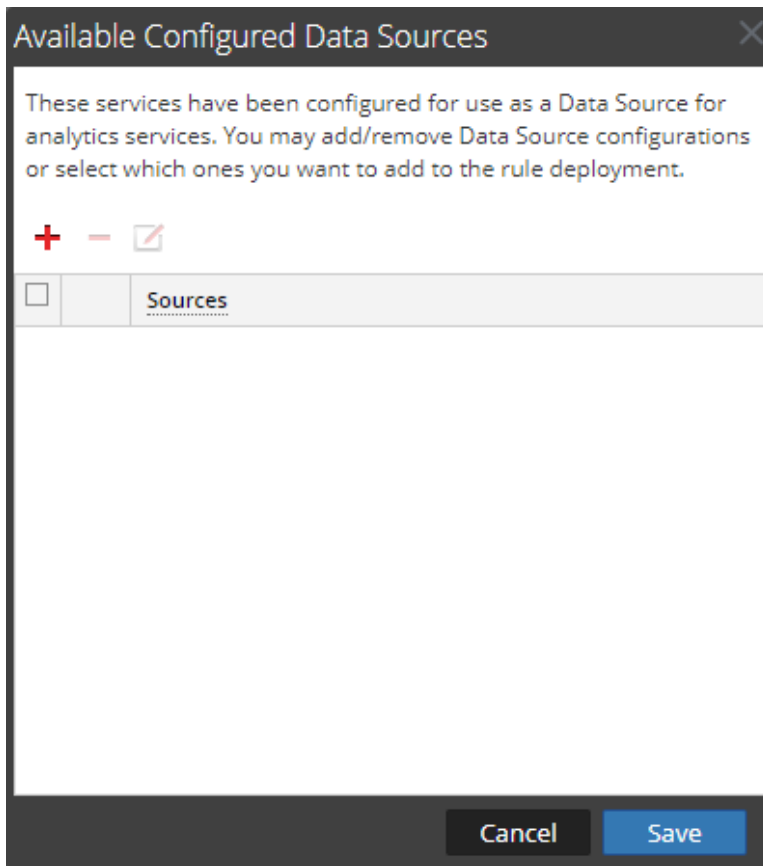
4. Select an ESA service and click **Save**.
The Deployment view is displayed. The ESA service is listed in the **ESA Services** section, with the status **Added**.

Step 3. Add Data Sources


Note: This option is available in version 11.3 and later.


You can select one or more data sources, such as Concentrators, to use for your selected ESA Service. This enables you to specify different data sources for each deployment. For example, you may want to use Concentrators with HTTP packet data in one deployment and Concentrators with HTTP log data in another deployment.

1. Go to  (Configure) > ESA Rules > Rules tab.
2. In the options panel, select a deployment.
3. Configure one or more data sources for your deployment. Do the following for each data source:
 - a. In the **Deployment** view **Data Sources** section, click  .
The **Available Configured Data Sources** dialog lists the services that have been configured for use as a data source.



- b. To add a data source configuration, click **+**.

The **Available Services** dialog lists the available data sources from the  (Admin) > Services view, such as Concentrators.



Note: You can add a Log Decoder as a data source for ESA, but it is better to add a Concentrator to take advantage of undivided aggregation as the Decoder may have other processes aggregating from it.

- c. In the **Available Services** dialog, select a data source, such as a Concentrator, and click **OK**.
- d. In the **Add Service** dialog, type the Administrator username and password for the data source.

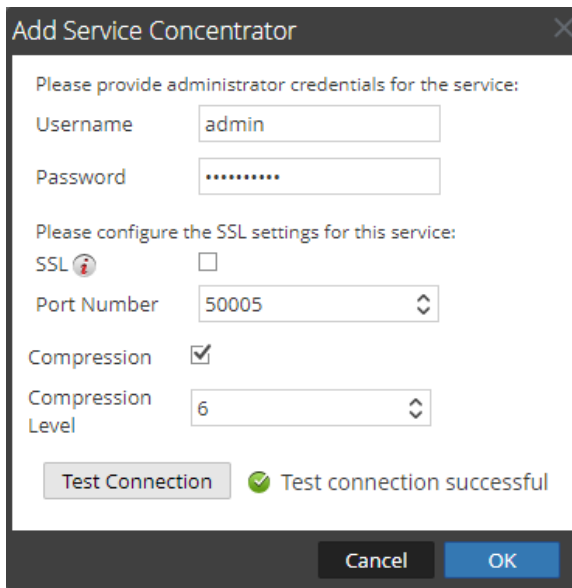
- e. To enable the SSL or Compression options, select the corresponding checkboxes.
- f. (Optional) You have the option to adjust the Compression Level for Concentrators on ESA in NetWitness 11.3 and later. To enable compression, select the **Compression** checkbox. You can set the **Compression Level** for a Concentrator from 0-9:

- Compression Level = **0** (If compression is enabled, it allows Core Services to control the amount of compression.)
- Compression Level = **1** (It uses the lowest amount of compression and has the highest performance.)
- Compression Level = **9** (It uses the highest amount of compression and has the worst performance.)

Somewhere in the middle between 1 and 9 is usually the best setting, which is what you get when you select a compression level of 0. For more detailed information, see the *Core Database Tuning Guide*.

Note: If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source. When you set the compression level for a Concentrator on ESA, it sets the same compression level for that Concentrator for ESA Correlation Rules.

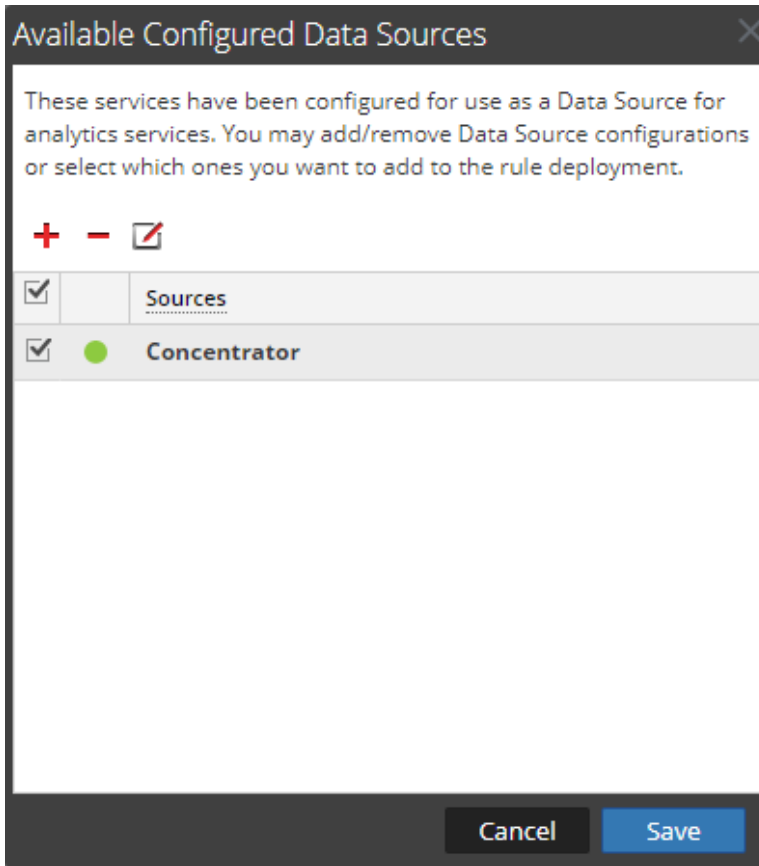
- g. Click **Test Connection** to make sure that it can communicate with the ESA service.



- h. Click **OK**.

After you configure your data sources and they appear in the **Available Configured Data Sources** dialog, you can use them for your deployment.

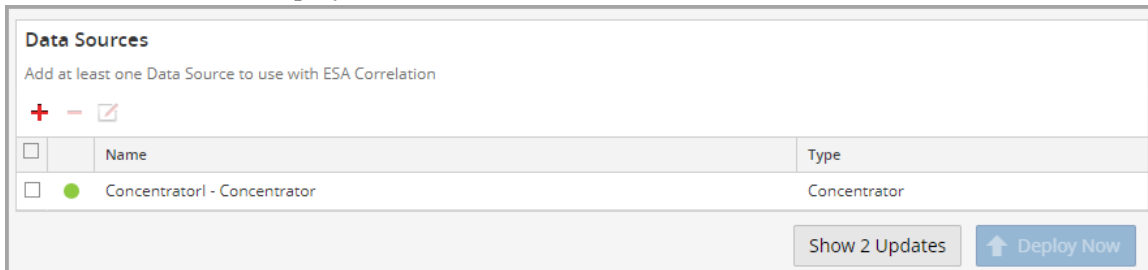
- 4. In the **Available Configured Data Sources** dialog, select at least one data source to use for the deployment.



A solid colored green circle indicates a running service and a white circle indicates a stopped service.

5. Click **Save**.

In the Deployment view **Data Sources** section, the selected data sources are added to the deployment. The **Deploy Now** button activates after an ESA service, a data source, and rules are added to an ESA rule deployment.





6. (Optional) If you have a medium to large NetWitness Platform deployment and you have high throughputs, you can create a filter to forward only the data relevant to this deployment to ESA. Before deploying the ESA rule deployment, add the data source filter. See [\(Optional\) Add a Data Source Filter](#).

Step 4. Add and Deploy Rules


This topic explains how to add ESA rules to an ESA rule deployment and then deploy the rules on ESA. Each ESA rule has unique criteria. The ESA rules in an ESA rule deployment determine which events ESA captures, which in turn determine the alerts you receive.

For example, Deployment A includes ESA Paris and, among others, a rule to detect file transfer using a non-standard port. When ESA Paris detects a file transfer that matches the rule criteria, it captures the event and generates an alert for it. If you remove this rule from Deployment A, ESA will no longer generate an alert for such an occurrence.

To add and deploy rules:

1. Go to  (Configure) > ESA Rules > Rules tab.
2. In the options panel, select a deployment.
3. In the **Deployment** view, click  in **ESA Rules**.

The Deploy ESA Rules dialog is displayed and shows each rule in your Rule Library:



<input type="checkbox"/>	Rule Name ^	Description	Trial Rule	Type
<input type="checkbox"/>	Endpoint Risk Scoring Rule Bundle	This bundle contains many rules which are used for ris...	No	Endpoint Rule B...
<input type="checkbox"/>	SAMPLE - Blacklist - From inside countries that are not ...	Monitors for non-SMTP traffic on TCP destination port ...	Yes	Rule Builder
<input type="checkbox"/>	SAMPLE - Non SMTP Traffic on TCP Port 25 Containing ...	Monitors for non-SMTP traffic on TCP destination port ...	Yes	Rule Builder
<input type="checkbox"/>	SAMPLE - P2P Software as Detected by an Intrusion De...	P2P software as detected by an intrusion detection dev...	Yes	Rule Builder
<input type="checkbox"/>	SAMPLE - User Added to Admin Group Same User su s...	Alert when user is upgraded to one of admin groups (c...	Yes	Rule Builder
<input type="checkbox"/>	SAMPLE - Whitelist - From outside of Germany, P2P Sof...	Whitelist Germany from P2P software as detected by a...	Yes	Rule Builder


4. Select rules and click **Save**.
The Deployment view is displayed and the **Deploy Now** button is enabled.


The screenshot displays the RSA NetWitness Platform interface for configuring a deployment. The main content area is titled "Deployment - Sample" and includes instructions: "Deployments map rules from your rule library to the appropriate ESA Services. Choose Rules, Services and rule execution method." Below this are three sections:

- ESA Services:** A table with columns: Status, Name, Address, Version, Last Deployment Date. One row is shown: "ESA Correlation" with status "Added" (indicated by a red exclamation mark icon), address "10.10.10.10", and version "11.5.0.0".
- Data Sources:** A table with columns: Name, Type. One row is shown: "Concentrator" with type "Concentrator".
- ESA Rules:** A table with columns: Status, Rule Name, Trial Rule, Severity, Type, Email, Snmp, Syslog, Script, Last Modified. Five rows are shown, all with status "Added".

At the bottom of the main content area, there is a "Deploy Now" button. The left sidebar shows "DEPLOYMENTS" with a "Sample" deployment listed, accompanied by a red exclamation mark icon and a "5" in a circle, indicating the number of rules in the deployment.

5. The rules are listed in the ESA Rules section.

- In the Status column, **Added** is next to each new rule.
- In the Deployments section,  indicates there are updates to the deployment.
- The total number of rules in the deployment is on the right.



6. Click **Deploy Now**.

The ESA Correlation service runs the rule set. After ESA correlation completes the processing of each rule in the deployment, the status changes to **Deployed**.

(Optional) Add a Data Source Filter

Note: This option is available in NetWitness Platform version 11.5 and later.

To improve performance, you can add an optional data source filter to your ESA rule deployment so that only the data relevant to the deployment is forwarded to ESA. The filter is comprised of application rules, which are applied to the Decoders mapped to your selected data sources. There are two options that you can use to create the filter:



- The **Simple** data source filter option enables you to select the application rules to be included in the filter query. The application rules that you select must be enabled on the Decoders that feed the data sources in your ESA rule deployment. ESA Correlation uses the filtered event data to process the ESA rules. This is the procedure described here.
- The **Advanced** data source filter option enables you to add a data source query directly. The individual application rule queries must be separated by an “or” condition. For more information on creating and writing Decoder rules, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.


Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.


IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

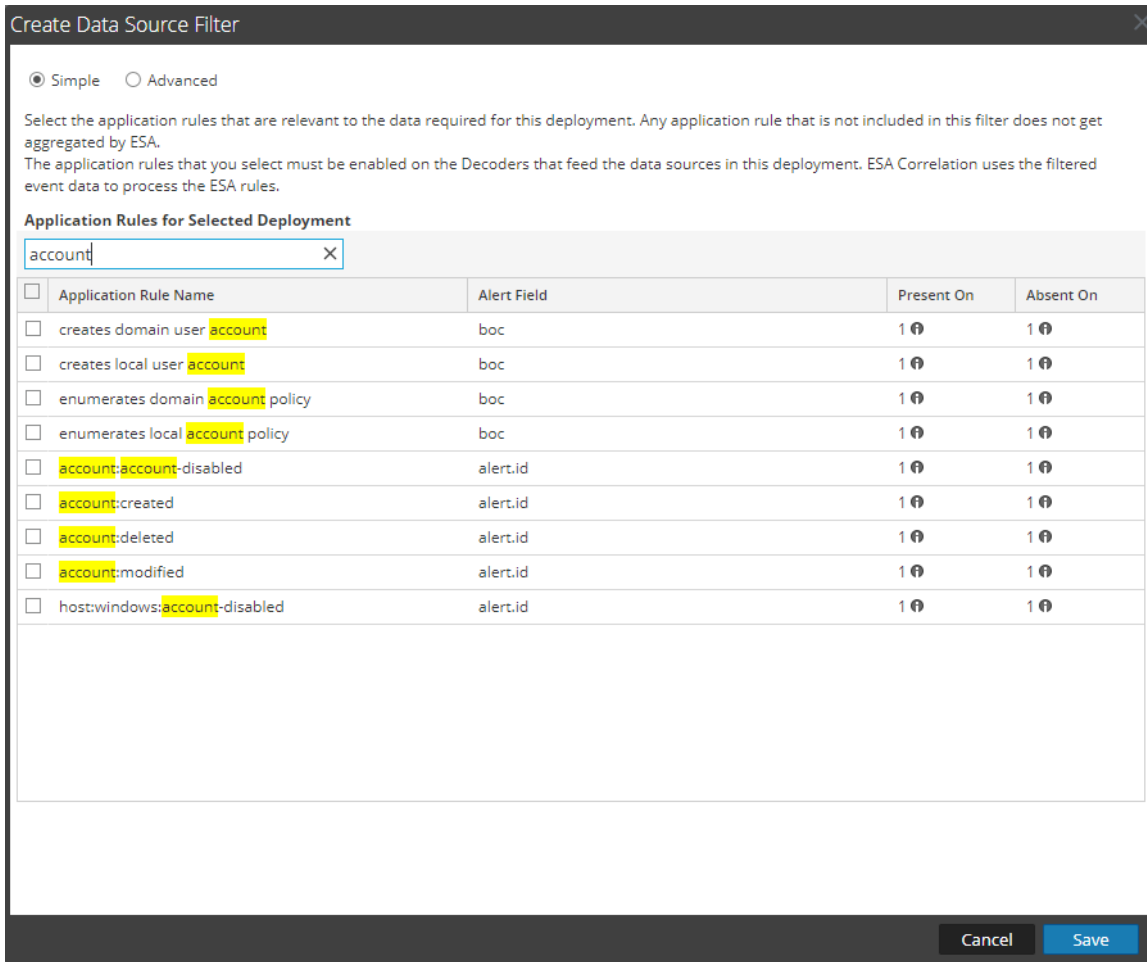
To add a Simple data source filter:

1. Go to  (Configure) > ESA Rules > Rules tab.
2. In the options panel, select a deployment.
3. In the **Deployment** view **Data Source Filter (Optional)** section, click  .




The **Create Data Source Filter** dialog lists the application rules that are available to filter the events in the data sources in your ESA rule deployment.

4. Select the application rules that you want to use to filter the data sources in your deployment and click **Save**. The **Alert Field** shows the meta key used in the alert. **Present On** shows the number of Decoders mapped to the data sources that have the rule. **Absent On** shows the number of mapped Decoders that do not have the rule. If present, hover over the  icon to view the names of the Decoders. You can use the filter to help locate the rules. For example, you can type "account" to search for application rules that contain that word.



The application rules that you select appear in a Filter Query with the status of Added.



- When you are ready to deploy your ESA rule deployment, click **Deploy Now**.

The ESA Correlation service runs the rule set. After ESA Correlation completes the processing of each rule in the deployment, the status changes to **Deployed**. The data source filter status changes to **Deployed** when the filter is actively streaming only the relevant data as defined in the filter query.

Troubleshooting Information: When filtering out a large portion of the traffic, you may see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file. Decrease the `max-sessions` parameter from 10,000 to a lower session count. See [Adjust Maximum Sessions for the Data Source Filter](#).

Deploy the Endpoint Risk Scoring Rules Bundle

An Endpoint Risk Scoring Rules Bundle, which contains approximately 400 rules, comes with NetWitness 11.3 and later. Endpoint risk scoring rules only apply to NetWitness Endpoint. You can add the Endpoint Risk Scoring Rules Bundle to an ESA rule deployment in the same way that you would add any ESA rule. However, you must specify endpoint data sources (Concentrators) in the ESA Rule Deployment.

The ESA Correlation service can process endpoint risk scoring rules, which generate alerts that are used in risk scoring calculations to identify suspicious files and hosts. To turn on risk scoring for NetWitness Endpoint, you must deploy endpoint risk scoring rules on ESA. For instructions, see "Deploy Endpoint Risk Scoring Rules on ESA" in the *ESA Configuration Guide*. For complete information on configuring NetWitness Endpoint, see the *NetWitness Endpoint Configuration Guide*.

Go to the [NetWitness Documents - All Versions](#) to find all RSA NetWitness Platform 11.x documents.


Additional ESA Rule Deployment Procedures

In addition to deploying an ESA service and rules, you may want to perform other steps on your ESA rule deployment, such as replacing an ESA service, changing a data source, editing or deleting a rule from the deployment, renaming or deleting the deployment, or showing updates to an ESA rule deployment.

Note: You cannot edit or duplicate an Endpoint Risk Scoring Rules Bundle.

In NetWitness version 11.3 and later, you can add or remove a data source from a deployment. In NetWitness 11.3.0.2 and later, you can edit a data source in an ESA rule deployment. This enables you to change the data source password, SSL, port, and compression settings.

- [Replace an ESA Service in an ESA Rule Deployment](#)
- [Edit a Data Source in an ESA Rule Deployment](#) (This option is available in NetWitness version 11.3.0.2 and later.)
- [Add or Remove a Data Source](#) (This option is available in version 11.3 and later.)
- [\(Optional\) Add or Remove a Data Source Filter](#) (This option is available in version 11.5 and later)
- [Edit or Delete a Rule in a Deployment](#)
- Stop Deployment
- [Edit the ESA Rule Deployment Name or Delete a Deployment](#)
- [Show Updates to an ESA Rule Deployment](#)
- [Fast Deployment and Full Deployment of ESA Rules](#)



Each of the following procedures starts in the Rules tab [ (Configure) > ESA Rules > Rules tab].

Anytime you make changes to an ESA rule deployment, you must redeploy it for the changes to take effect. To redeploy the deployment, click the **Deploy Now** button for that deployment.

Replace an ESA Service in an ESA Rule Deployment

An ESA rule deployment can have only one ESA service, but you can replace it at any time with another ESA service. You can use the same ESA service in multiple deployments.

Remove an ESA Service from an ESA Rule Deployment

1. Go to  (Configure) > ESA Rules > Rules tab.
The Rules tab is displayed.
2. In the options panel, under **Deployments**, select a deployment.
3. In the **ESA Services** section, select a service and click  in the toolbar.
A confirmation dialog is displayed.

4. Click **Yes**.
The service is removed from the deployment.
5. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Add an ESA Service to an ESA Rule Deployment



1. To add an ESA Service to an ESA rule deployment, see [Step 2. Add an ESA Service](#). For the ESA Correlation service in NetWitness Respond 11.3 and later, you must add at least one data source to the service. See [Step 3. Add Data Sources](#).
2. After you finish making changes to the ESA rule deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the ESA rule deployment is redeployed.

Edit a Data Source in an ESA Rule Deployment

Note: This procedure applies to NetWitness Platform 11.3.0.2 and later versions.

You can change the configuration of a data source in an ESA rule deployment. You can change the data source password, SSL, port, and compression settings. When a data source password changes, it is important to change the password on the data source so that ESA can continue to communicate with the data source.

Note: If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source.

1. Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **Data Sources** section, select a data source and click  in the toolbar.

- In the **Edit Service** dialog, type the Administrator username and password for the data source. If the password changed on the data source, enter the new password here.

Edit Service LH - Concentrator

Please provide administrator credentials for the service:

Username:

Password:

Please configure the SSL settings for this service:

SSL

Port Number:

Compression

Compression Level:

- To enable the SSL or Compression options, select the corresponding checkboxes.
- (Optional) You have the option to adjust the Compression Level for Concentrators on ESA in NetWitness 11.3 and later. To enable compression, select the **Compression** checkbox. You can set the **Compression Level** for a Concentrator from 0-9:
 - Compression Level = **0** (If compression is enabled, it allows Core Services to control the amount of compression.)
 - Compression Level = **1** (It uses the lowest amount of compression and has the highest performance.)
 - Compression Level = **9** (It uses the highest amount of compression and has the worst performance.)

Somewhere in the middle between 1 and 9 is usually the best setting, which is what you get when you select a compression level of 0. For more detailed information, see the *Core Database Tuning Guide*.

Note: When you set the compression level for a Concentrator on ESA, it sets the same compression level for that Concentrator for ESA Correlation Rules.



- Click **Test Connection** to make sure that it can communicate with the ESA service.

- Click **OK**.
- After you finish making changes to the deployment, click **Deploy Now** to redeploy the ESA rule deployment. The changes take effect on ESA after the deployment is redeployed. You can view the update information in the Updates to the Deployments dialog. See [Show Updates to an ESA Rule Deployment](#).

Add or Remove a Data Source

Note: This option is available in NetWitness version 11.3 and later.

Remove a Data Source from an ESA Rule Deployment

- Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.
- In the Rules tab options panel, under **Deployments**, select a deployment.
- In the **Data Sources** section, select the data source and click  in the toolbar.
The data source is removed from the deployment.
- After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Add a Data Source to an ESA Rule Deployment

- To add a data source, see [Step 3. Add Data Sources](#).
- After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

(Optional) Add or Remove a Data Source Filter

To improve performance, you can add an optional data source filter to your ESA rule deployment so that only the data relevant to the deployment is forwarded to ESA. The filter is comprised of application rules, which are applied to the Decoders mapped to your selected data sources.

You cannot edit a data source filter. To modify a data source filter, you must remove the filter, add a new filter, and then redeploy the ESA rule deployment.

Note: This option is available in NetWitness version 11.5 and later.



Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.

Remove a Data Source Filter from an ESA Rule Deployment

If you plan to replace your data source filter with an adjusted filter, you may want to copy the filter query in the data source filter before you remove it and compare it with the new query.

1. Go to  **(Configure)** > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **Data Source Filter (Optional)** section, select the filter and click  in the toolbar.
The filter is removed from the deployment.
4. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Add a Data Source Filter to an ESA Rule Deployment

1. To add a simple data source filter, see [\(Optional\) Add a Data Source Filter](#). If necessary, you can use the advanced filter instead of the simple filter to add a data source query directly. The individual application rule queries must be separated by an "or" condition. For more information on creating and writing Decoder rules, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.
2. After you finish making changes to the deployment, click **Deploy Now** to redeploy it. The changes take effect on ESA after the deployment is redeployed.

Adjust Maximum Sessions for the Data Source Filter



When filtering out a large portion of the traffic, you may see an "Invalid header size" error while communicating with Core services in the ESA Correlation log file. (You can use SSH to get in the system and go to: `/var/log/netwitness/correlation-server/correlation-server.log`).

Lower the `max-sessions` parameter until you no longer see the error in the log. The more you filter out the traffic, the lower you should set the `max-sessions` parameter.

1. In the Explore view node list for an ESA Correlation service, select **correlation > stream**.

The screenshot shows the RSA NetWitness Platform interface. The left sidebar shows a tree view of services, with 'correlation' and 'stream' highlighted. The main panel displays the configuration for the 'stream' service. The 'max-sessions' parameter is highlighted with a red box and has a value of 10000.


Parameter	Value
event-source-id	true
filter	
idle-retry-interval	10
lag-time	15 MINUTES
lowercase	
max-sessions	10000
mechanism	AGGREGATION
metrics	
minutes-back	5
multi-valued	action , alert , alertid , alias.host , alias.ip , alias.ipv6 , analysis.file , analysis.service , analysis.session , boc , browserprint , cert.thumbprint , checksum , checksum.all , checksum.dst , checksum.src , client.all , content , context , context.all , context.dst , context.src , dir.path , dir.path.dst , dir.path.src , directory , directory.all , directory.dst , directory.src , email , email.dst , email.src , eoc , feed.category , feed-desc , feed-name , file.cat , file.cat.dst , file.cat.src , filename.dst , filename.src , filter , function , host.all , host.dst , host.orig , host.src , host.state , inv.category , inv.context , loc , ip.orig , ipv6.orig , netname , OS , param , param.dst , param.src , registry.key , registry.value , risk , risk.info , risksuspicious , risk.warning , threat.category , threat.desc , threat.source , user.agent , username
multi-valued-as-array	false
no-system-meta	false
pre-fetch	2
query	
reader-buffer-size	1048576
retrieve-record-stream-stats-every	2 MINUTES
retrieve-schema-every	5 MINUTES
retrieve-schema-frequency	1
retry-timeout	10000
save-position-every	1 MINUTES
save-position-frequency	1
single-valued	accesses , context.target , file.attributes , logon.type.desc , packets
socket-timeout	0

2. In `max-sessions`, lower the value until you no longer see the error in the ESA Correlation log file. The default value is **10000**.
3. Restart the ESA Correlation service. Go to  (Admin) > Services, select the ESA Correlation service, and then select  > Restart.

Edit or Delete a Rule in a Deployment



In an ESA rule deployment, you can edit and delete rules to customize the deployment.

Edit a Rule



1. Go to  (Configure) > ESA Rules > Rules tab. The Rules tab is displayed.
2. In the Rules tab options panel, under **Deployments**, select a deployment.
3. In the **ESA Rules** panel, double-click a rule to open it in a new tab.

4. Modify the rule, then click **Save**.
The rule is saved.
5. Click **Deploy Now** to redeploy the deployment.
The changes take effect on ESA after the deployment is redeployed.

Delete a Rule


1. Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.
2. In the options panel, under **Deployments**, select a deployment.
3. In the **ESA Rules** panel, select a rule and click  in the toolbar.
A confirmation dialog is displayed.
4. Click **Yes**.
The rule is deleted.
5. Click **Deploy Now** to redeploy the deployment.
The changes take effect on ESA after the deployment is redeployed.

Stop Deployment



1. Go to  (**Configure**) > **ESA Rules** > **Rules** tab.
The Rules tab is displayed.
2. In the options panel, under **Deployments**, select a deployment.
3. In the **ESA Rules** panel, select a rule and click  in the toolbar.
A confirmation dialog is displayed.
4. Click **Yes**.
The rule is to stopped.
5. Click **Stop Deployment** to stop the deployment.

Edit the ESA Rule Deployment Name or Delete a Deployment


To access the deployments:

1. Go to  (**Configure**) > **ESA Rules**.
The Configure view is displayed with the Rules tab open.
2. In the options panel, under **Deployments**, select a deployment.
The Deployment view is displayed.

Edit the ESA Rule Deployment Name

1. In the options panel, under **Deployments**, select a deployment.
The Deployment view is displayed.
2. Select  > **Edit**.
The deployment name is made available for editing.
3. Enter the new deployment name.
4. Click **Deploy Now** to redeploy the deployment.
The changes take effect on ESA after the ESA rule deployment is redeployed. In NetWitness 11.3 and later, the deployment names that you choose appear on the deployment tabs in the  **(Configure) > ESA Rules > Services** tab.


Delete an ESA Rule Deployment

1. In the options panel, under **Deployments**, select a deployment.
The Deployment view is displayed.
2. Select  > **Delete**.
A confirmation dialog is displayed.

3. Click **Yes**.


The deployment is deleted.

Show Updates to an ESA Rule Deployment

You can view changes to an ESA rule deployment, such as adding or removing rules. When there is a change to a deployment, the update icon () appears next to the name of the deployment in the Rules tab options panel.

1. Go to  (**Configure**) > **ESA Rules**.

The Rules tab is displayed.

2. In the options panel, under **Deployments** click **Show Updates** on the far right.


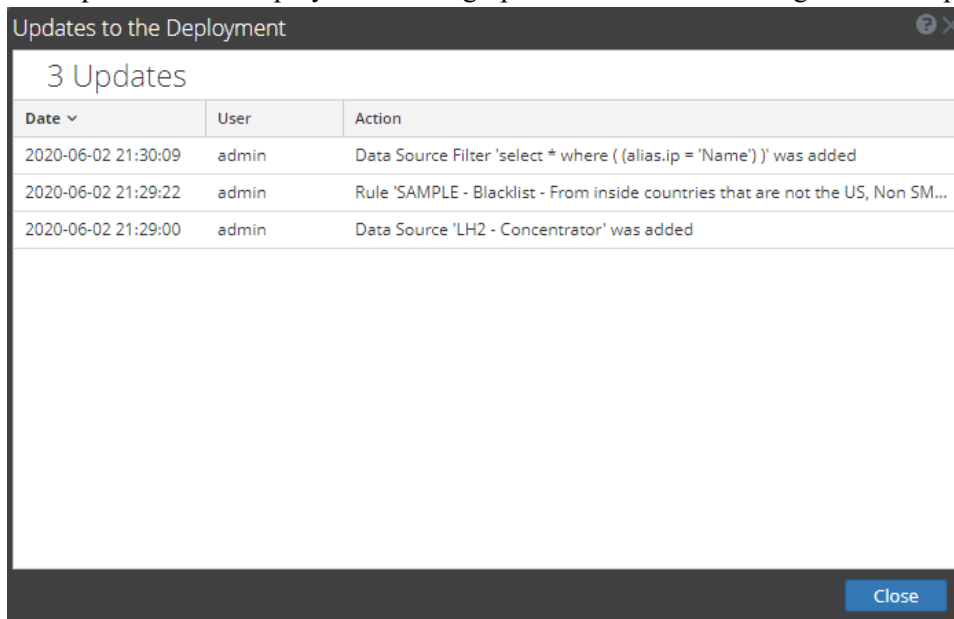
The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation bar has 'LIVE CONTENT', 'SUBSCRIPTIONS', 'INCIDENT RULES', 'INCIDENT NOTIFICATIONS', 'ESA RULES', 'CUSTOM FEEDS', 'EVENT RULES', and 'LOG PARSER RULES'. The 'Rules' tab is selected, and the 'Services' and 'Settings' sub-tabs are visible. The 'DEPLOYMENTS' section on the left shows 'Deployment A' and 'Deployment B'. The 'ESA Rules' section displays a table of rules with the following data:

Status	Rule Name	Trial Rule	Severity	Type	Email	Snmp	Syslog	Script	Last Modified
Deployed	SAMPLE - Non SMTP Traffic on TCP Port 25 Cont...	Yes	Low	Rule Builder	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	2019-07-26 05:49:21
Deployed	SAMPLE - P2P Software as Detected by an Intrus...	Yes	Low	Rule Builder	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	2019-07-26 05:49:21
Deployed	SAMPLE - User Added to Admin Group Same Us...	Yes	Medium	Rule Builder	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	2019-07-26 05:49:21

Below the table, there is a 'Data Source Filter (Optional)' section with a table showing filter queries:

Status	Filter Query	Last Modified
Added	select * where ((alias.ip = 'Name'))	2020-06-02 17:30:09

The Updates to the Deployments dialog opens and shows the changes to the deployment.



3. Click **Close**.


Fast Deployment and Full Deployment of ESA Rules

Only changes to an ESA rule is eligible for a fast deployment. This means if an ESA rule has been added, modified or deleted only then you can perform a fast deployment. In case, along with rules, if anything else has been changed such as a data source has been added, then the Fast Deployment option will not be displayed as it will not be eligible. The Fast Deployment option can be used to quickly load the latest rule that runs entirely on the Esper Engine hence reducing the need to redeploy the engine again in a shorter time where all the related settings of the selected deployment are deleted and the deployment is recreated from scratch.

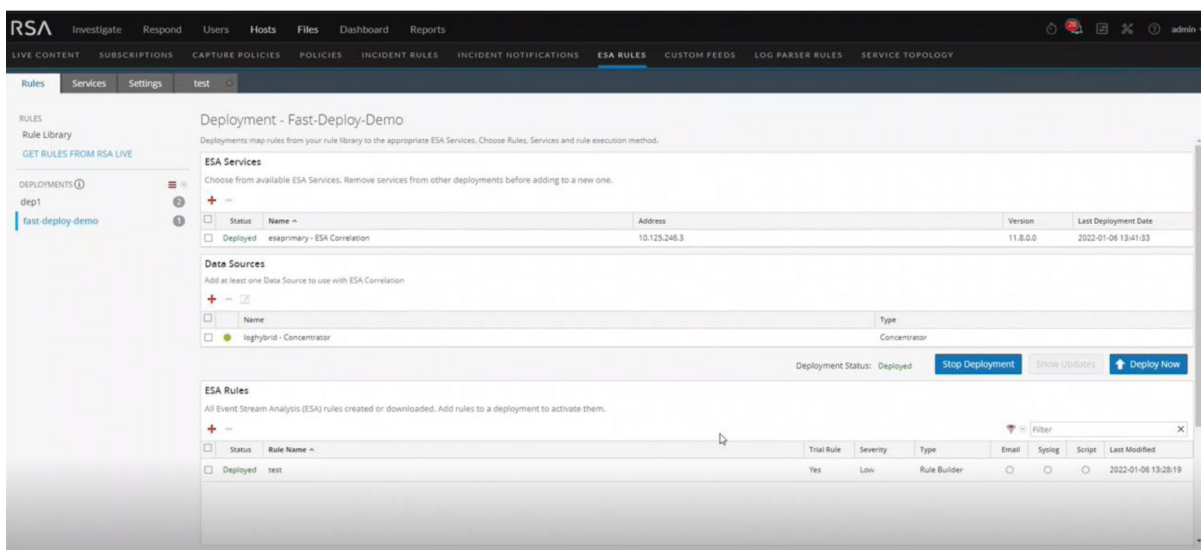
The Full Re-Deploy can be used if:


- The previously deployed rules failed, and the analyst wants to redeploy the same rules again.
- There are other changes made to the ESA Services, the Data Sources along with the ESA Rules.
- If the analyst wants to tear down the entire engine and recreate it.

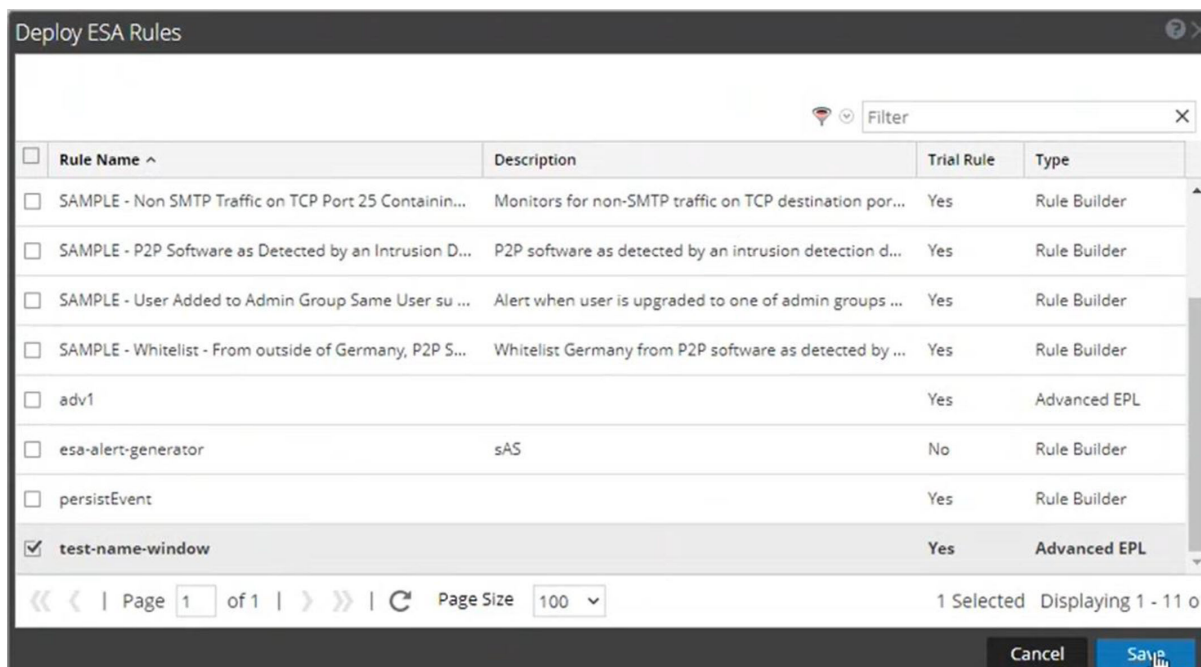
To perform Fast Deployment:




1. Go to **(Configure) > ESA Rules**.
The Rules tab is displayed.



2. Under ESA Rules, click  to deploy a new rule.
3. Select a rule you want to deploy from the list and click **Save**.



4. Click **Deploy Now**.
A confirmation message is displayed.



5. Click **Fast Deploy** to begin the deployment of the new ESA Rule.

Note: The Fast Deploy button will be displayed only if the ESA rule has been added, deleted or updated, not if you have made updated to the ESA Services or Data Sources.

View ESA Stats and Alerts

When ESA generates alerts, you can view details about how the rules performed, such as statistics on the engine, rule, and alert, and you can also view information on which rules are enabled or disabled. For instructions on viewing ESA stats, see [View Stats for an ESA Service](#)


When your ESA generates alerts, you can view the results in the Respond Alerts List view. This enables you to see trends and understand both the volume and frequency of alerts. For instructions on viewing alerts, see [View a Summary of Alerts](#)


View Stats for an ESA Service

This topic describes how to view the deployment statistics (stats) for an ESA Correlation service. This procedure is useful when you are attempting to determine the effectiveness of a rule or troubleshoot an ESA rule deployment.

Caution: When you modify and re-deploy an ESA rule deployment, all of the stats are removed from that deployment. The generated alerts are not removed from NetWitness Respond.

View ESA Stats

1. Go to  (Configure) > ESA Rules > Services tab.
2. From the ESA Services list on the left, select a service.
The deployment stats for the selected service are displayed.



The screenshot shows the RSA NetWitness Platform interface for the 'esaSecondary - ESA Correlation' service. The page is divided into several sections:

- Engine Stats:** Esper Version: 8.4.0, Time: 2020-05-05T19:44:43, Events Offered: 74750, Offered Rate: 0 per second / 8,557 max, Status: Active.
- Rule Stats:** Rules Enabled: 75, Rules Disabled: 0, Events Matched: 29.
- Alert Stats:** Notifications: 0, Message Bus: 1.
- Deployed Rule Stats:** A table listing various rules with columns for Enabled, Name, Rule Type, Trial Rule, Last Detected, Events Matched, Memory Usage, and CPU %.

Enabled	Name	Rule Type	Trial Rule	Last Detected	Events Matched	Memory Usage	CPU %
<input type="checkbox"/>	AWS Critical VM Modified	Esper	Yes	0	0	0 bytes	0.017
<input type="checkbox"/>	AWS Permissions Modified Followed By Instan...	Esper	Yes	0	0	168 bytes	0.069
<input type="checkbox"/>	Account Added to Administrators Group and ...	Esper	Yes	0	0	168 bytes	0.134
<input type="checkbox"/>	Account Removals From Protected Groups on ...	Esper	Yes	0	0	64 bytes	0.047
<input type="checkbox"/>	Aggressive Internal Database Scan	Esper	Yes	0	0	64 bytes	0.087
<input type="checkbox"/>	Aggressive Internal NetBIOS Scan	Esper	Yes	0	0	4.22 KB	0.851
<input type="checkbox"/>	Aggressive Internal Web Portal Scan	Esper	Yes	0	0	1016 bytes	0.253
<input type="checkbox"/>	BYOD Mobile Web Agent Detected	Esper	Yes	0	0	64 bytes	0.079
<input type="checkbox"/>	Backdoor Activity Detected	Esper	Yes	0	0	0 bytes	0.014
<input type="checkbox"/>	Cerber Ransomware	Esper	Yes	0	0	896 bytes	0.314


- (This option applies to NetWitness version 11.3 and later.) In the Deployment view under the ESA Correlation service name, select the tab of the deployment you would like to view. For example, select the Deployment A tab to view the stats for deployment A. Select the Deployment B tab to view the status for deployment B.
- Review the following sections of ESA stats.
For a complete description of each statistic in each section, see [Services Tab](#).
 - Engine Stats**
 - Rule Stats**
 - Alert Stats**
- In the **Deployed Rule Stats**, review details about the rules deployed on the ESA.
For a complete description of each column in each section, see [Services Tab](#).
 - If the rule is enabled or disabled
 - What the rule name is
 - The type of rule
 - If the rule is running in Trial Rule mode
 - Last detected
 - Events matched
 - The amount of memory used by the rule
 - The percentage of the deployment CPU used by the rule (available in NetWitness version 11.5 and later)
- To monitor overall memory usage and health of your ESA Correlation service, click **Health & Wellness**.

Enable or Disable Rules

- In the **Deployed Rule Stats** panel, select a rule from the grid.
- Click **Enable** to enable the rule, or click **Disable** to disable the rule.
The Services tab is refreshed to show the changes, which take effect immediately.

Refresh the Statistics

The Services tab does not update statistics automatically unless you enable or disable a rule. To ensure you view current statistics:

- Click  in the upper right corner to refresh the information.
- View the updated information.

View a Summary of Alerts

In the Repond view, you can browse through various alerts from multiple sources. You can filter the alerts list to show only alerts of interest, such as by Alert Name, alert source, and a specific time range from the following sources:

1. Detect AI
2. Endpoint
3. Event Stream Analysis
4. Malware Analysis
5. NetWitness Investigate
6. Reporting Engine
7. Risk Scoring
8. User Entity Behavior Analysis
9. Web Threat Detection

Perform the following steps to use the functionalities provided in the Respond view.

1. Go to **Respond > Alerts**.

The Respond Alerts List view displays a list of all NetWitness alerts.

Created	Severity	Name	Source	# Events	Host Summary	Incident ID
06/02/2020 10:30:24 pm	10	IP Source is 10.1633 to 10.1633	Event Stream Analysis	1	10.1633 to 10.1633	INC-1070
06/02/2020 10:30:24 pm	50	IP Source is 10.61949 to 10.61949	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:30:24 pm	20	Alert without Incident	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:29:36 pm	20	IP Source is 192.192	Event Stream Analysis	1	192.192	INC-1060
06/02/2020 10:29:19 pm	10	IP Source is 10.1633 to 10.1633	Event Stream Analysis	1	10.1633 to 10.1633	INC-1062
06/02/2020 10:29:19 pm	50	IP Source is 10.61949 to 10.61949	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:29:19 pm	20	Alert without Incident	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:28:14 pm	10	IP Source is 10.1633 to 10.1633	Event Stream Analysis	1	10.1633 to 10.1633	INC-1068
06/02/2020 10:28:14 pm	20	Alert without Incident	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:27:32 pm	20	IP Source is 192.192	Event Stream Analysis	1	192.192	INC-1060
06/02/2020 10:27:10 pm	50	IP Source is 10.61949 to 10.61949	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:27:10 pm	20	Alert without Incident	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:27:10 pm	10	IP Source is 10.1633 to 10.1633	Event Stream Analysis	1	10.1633 to 10.1633	INC-1067
06/02/2020 10:26:34 pm	90	Threshold Breached for FILE_SMB_SyncHostLoss	Risk Scoring	1		INC-23
06/02/2020 10:26:34 pm	90	Threshold Breached for FILE_ExecHostLoss	Risk Scoring	1		INC-23
06/02/2020 10:26:34 pm	90	Threshold Breached for FILE_SMB_SyncHostLoss	Risk Scoring	1		INC-22
06/02/2020 10:26:34 pm	90	Threshold Breached for FILE_ExecHostLoss	Risk Scoring	1		INC-21
06/02/2020 10:26:08 pm	20	Alert without Incident	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:26:08 pm	10	IP Source is 10.1633 to 10.1633	Event Stream Analysis	1	10.1633 to 10.1633	INC-1066
06/02/2020 10:26:08 pm	50	IP Source is 10.61949 to 10.61949	Event Stream Analysis	1	10.61949 to 10.61949	INC-1039
06/02/2020 10:25:38 pm	90	Blacklisted File	Endpoint	1		

2. In the **Filters** panel on the left, you can filter the alerts list to view specific alerts for a specific time frame. For example, in the Alert Names section, you can select an alert for an ESA rule, such as

Direct Login to an Administrative Account, and leave the Time Frame set to Last Hour. The alerts list to the right shows a list of alerts that match your filter selection along with a count of the alerts at the bottom of the alerts list.

The alerts list shows information about each of the alerts.

- **Created:** Displays the date and time when the alert was created in the source system.
 - **Severity:** Displays the level of severity of the alert. The values are from 1 to 100.
 - **Name:** Displays a basic description of the alert.
 - **Source:** Displays the original source of the alert.
 - **# of Events:** Indicates the number of events contained within an alert.
 - **Host Summary:** Displays details of the host, like the host name from where the alert was triggered.
 - **Incident ID:** Shows the incident ID of the alert. If there is no incident ID, the alert does not belong to an incident.
3. You can click an alert in the list to open an **Overview** panel on the right where you can view raw alert metadata.

The screenshot displays the RSA NetWitness Respond interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboards', and 'Reports'. The main area is divided into 'INCIDENTS', 'ALERTS', and 'TASKS'. The 'ALERTS' tab is active, showing a list of alerts with columns for 'CREATED', 'SEVERITY', 'NAME', 'SOURCE', '# EVENTS', 'HOST SUMMARY', and 'INCIDENT ID'. The first alert is highlighted in blue.

CREATED	SEVERITY	NAME	SOURCE	# EVENTS	HOST SUMMARY	INCIDENT ID
06/02/2020 05:00:57 am	100	Direct Login To an Administrative Account	Event Stream Analysis	1	computer_...	INC-8
06/02/2020 05:00:57 am	96	abnormal_login_day_time	User Entity Behavior An...	7		
06/02/2020 10:26:34 pm	90	Threshold Reached for FILE_SMOvchost.exe	Risk Scoring	1		INC-25
06/02/2020 10:26:34 pm	90	Threshold Reached for FILE_vchost.exe	Risk Scoring	1		INC-24
06/02/2020 10:26:34 pm	90	Threshold Reached for FILE_SMOvchost.exe	Risk Scoring	1		INC-23
06/02/2020 10:26:34 pm	90	Threshold Reached for FILE_SMOvchost.exe	Risk Scoring	1		INC-22
06/02/2020 10:26:34 pm	90	Threshold Reached for FILE_vchost.exe	Risk Scoring	1		INC-21
06/02/2020 10:25:38 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 10:25:38 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 10:25:38 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 10:25:58 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 10:24:34 pm	90	Threshold Reached for FILE_cmd.exe	Risk Scoring	1		INC-17
06/02/2020 10:24:11 pm	90	Blacklisted File	Endpoint	1		INC-1064
06/02/2020 10:24:07 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 10:24:07 pm	90	Blacklisted File	Endpoint	1		INC-1064
06/02/2020 10:24:04 pm	90	Threshold Reached for FILE_...	Risk Scoring	1		INC-157
06/02/2020 10:24:02 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 10:23:56 pm	90	Blacklisted File	Endpoint	1		
06/02/2020 09:30:34 pm	90	Threshold Reached for FILE_SMOvchost.exe	Risk Scoring	1		INC-23
06/02/2020 09:30:34 pm	90	Threshold Reached for FILE_vchost.exe	Risk Scoring	1		INC-24
06/02/2020 09:30:34 pm	90	Threshold Reached for FILE_vchost.exe	Risk Scoring	1		INC-23

The right-hand pane shows the 'OVERVIEW' for the selected alert. It includes fields for Incident ID (INC-8), Created (06/02/2020 05:00:57 am), Severity (100), Source (Event Stream Analysis), Type (Log), and # Events (1). Below this is a 'Raw Alert' section containing a JSON object with various event details.

```

{
  "events": [
    {
      "checkus_src": "..._checkus_src",
      "reference_id": "4624",
      "agent_ip": "...",
      "device_src": "...",
      "device_type": "...",
      "event_source": "Microsoft-Windows-Security-Auditing",
      "extension": "2894",
      "id": "2894",
      "eth_src": "...",
      "event_type": "Information",
      "analysis_service": "..._analysis_service",
      "event_computer": "...",
      "logmic_src": "...",
      "action": "..._action",
      "msg_id": "..._Microsoft-Windows-Security-Auditing",
      "directory_src": "..._directory_src",
      "country_dst": "USA",
      "ec_subject": "User",
      "event_source_ip": "...",
      "msg_id": "148336870286",
      "top_report": "21",
      "checkus_dst": "..._checkus_dst",
      "mail_src": "..._mail_src",
      "ip_dst": "...",
      "device_ip": "...",
      "logmic_dst": "...",
      "event_desc": "An account was successfully logged on."
    }
  ]
}
    
```

For more information about filtering alerts and viewing alert details, see the *NetWitness Respond User Guide*.

Add an Advanced EPL Rule

This topic provides instructions to define rule criteria by writing an EPL query. EPL is a declarative language for handling high-frequency time-based event data. It is used to express filtering, aggregation, and joins over possibly sliding windows of multiple event streams. EPL also includes pattern semantics to express complex temporal causality among events.

Write an advanced EPL rule when rule criteria is more complex than what you can specify in Rule Builder.

It is outside the scope of this guide to explain EPL syntax.

- For EPL Documentation, see <http://www.espertech.com/esper/esper-documentation/>
- For the EPL Online Tool, see <http://esper-epl-tryout.appspot.com/epltryout/mainform.html>



For best practices on writing advanced EPL rules, see [ESA Rule Writing Best Practices](#).


Prerequisites

The following are prerequisites for adding an advanced rule:

- You must know Event Processing Language (EPL).
- You must understand ESA Annotations to mark which EPL statements are linked to generating alerts.

Add an Advanced EPL Rule


1. Go to  (Configure) > ESA Rules.
2. In the **Rule Library**, select  > **Advanced EPL**.



The screenshot shows the 'New Advanced EPL Rule' configuration page in the RSA NetWitness Platform. The page is titled 'Advanced EPL' and instructs the user to 'Write a rule in Event Processing Language.' The configuration fields are as follows:

- Rule Name ***: A text input field.
- Description**: A large text area for describing the rule.
- Trial Rule**: A checkbox that is checked.
- Memory Threshold**: A dropdown menu set to '100' and 'MB'.
- Alert**: A checkbox that is checked.
- Severity ***: A dropdown menu set to 'Low'.
- Query ***: A large text area for entering the EPL query.
- Notifications**: A section with a table for configuring notifications. The table has columns for 'Output', 'Notification', 'Notification Server', and 'Template'. Below the table, it says 'No parameters to edit.' There is also an option for 'Output Suppression of every' followed by a text input and 'minutes'.
- Enrichments**: A section with a table for configuring enrichments. The table has columns for 'Output', 'Enrichment Source', 'ESA Event Stream Meta', and 'Enrichment Source Column Name'. Below the table, it says 'No parameters to edit.'

The footer of the page shows 'RSA NETWITNESS PLATFORM' and the version '11.5.0.0'.

3. Type a unique, descriptive name in the **Rule Name** field.
This name will appear in the Rule Library so be specific enough to distinguish the rule from others.
4. In the **Description** field, explain which events the rule detects.
The beginning of this description will appear in the Rule Library
5. Select **Trial Rule** to automatically disable the rule if all trial rules collectively exceed the memory threshold.
Use trial rule mode as a safeguard to see if a rule runs efficiently and to prevent downtime caused by running out of memory. For more information, see [Working with Trial Rules](#).
6. (This option applies to 11.5 and later.) Enter a **Memory Threshold** for a rule that uses memory, such as a rule that contains windows or pattern matching. If the configured memory threshold is exceeded, the rule gets disabled individually and an error is displayed for that rule on the  (Configure) >

ESA Rules > Services tab. The Memory Threshold option works for trial rules and non-trial rules. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.

7. (This option applies to 11.3 and later.) Select **Alert** to send an alert to Respond. Clear the checkbox if you do not want to send an alert to Respond. To turn alerts on or off for ALL rules, see the *ESA Configuration Guide*.
8. For **Severity**, classify the rule as Low, Medium, High or Critical.
9. To define rule criteria, write a **Query** in EPL.

Note: For all meta key names, use an underscore not a period. For example, `ec_outcome` is correct but `ec.outcome` is not.

Supported meta entities:

S. No	Supported Meta Entities	Description
1	fullname_all	
2	eth_all	
3	ip_all	Combines all the IPv4 meta keys.
4	ipv6_all	Combines all the IPv6 meta keys.
5	port_src_all	
6	port_dst_all	
7	dir_path_all	
8	org_all	
9	geoip_all	
10	port_all	
11	domain_all	
12	email_all	
13	filename_all	
14	directory_all	
15	checksum_all	
16	param_all	
17	context_all	
18	attack_all	
19	analysis_all	

20	compromise_all	
21	inv_all	
22	outcome_all	
23	ec_all	
24	user_all	
25	host_all	
26	client_all	

10. For dynamic statement name generation in ESA, you must enclose the meta keys in curly brackets and include this annotation in the syntax:

```
@Name("RIG {ip_src} {alias_host} {ec_activity}")
```

where,

- RIG is the static part of the statement name
- {ip_src}, {alias_host}, {ec_activity} is the dynamic part of the statement name

Note: If any of the metas in the dynamic part of the statement name has a null value, it is displayed as a static text.

If a rule should generate an alert, include this ESA annotation in the syntax:

```
@RSAAlert
```

For more information on ESA Annotations, see [ESA Annotations](#).

Validate an Advanced EPL Rule

You can confirm that an ESA rule generates the expected alerts by testing the rule logic using JSON input data. You can view the alerts in the output, but this test does not send any alert notifications.

If you want to view all of the debug information, include an `@Audit('stream')` annotation to your rule query and view the Debug Log in the test output. To enable auditing you require to add `@Audit` to the rule and set the logging level for the Esper audit package to INFO. This can be done by creating `correlation-server.yml` file under `/etc/netwitness/correlation-server` with this content and restarting the `correlation-server` service with `systemctl restart rsa-nw-correlation-server`:


```
logging:
  level:
    com.espertech.esper.audit: INFO
```


The following basic example query contains the `@Audit('stream')` annotation and queries for events that do not have a source IP of 1.1.1.1 or 2.2.2.2.

```

Query *
@Audit('stream')
@RSAAAlert
SELECT * FROM Event((ip_src NOT IN ( '1.1.1.1' , '2.2.2.2' )));


```


1. If you are not already in the rule, go to  (**Configure**) > **ESA Rules** > **Rules** tab and in the **Rule Library**, open the ESA rule that you want to test.
2. Scroll down to the **Test Rule** section.





3. In the **ESA Service** field, select the ESA Correlation service to process the rule. Use the same ESA Correlation service that you plan to use in the ESA rule deployment that contains the rule.
4. In the **Input Data** field, enter the input events to test the rule. Download the events from the Investigate view in JSON format, copy the events, and paste them in this field. You can do this from the Investigate > Navigate view or the Investigate > Events view.

To download the events from the Investigate > Navigate view:

- a. In the main menu, go to **Investigate** > **Navigate** in a new tab, select a data source, and click **Navigate**.
- b. In the Navigate view, click **Load Values** and click a meta value to filter the events.
- c. Save the events as meta in the JSON file format [Save Events > Meta > (name the file) > Export Meta Format: choose JSON].
- d. In the toolbar click the  (**Jobs**) icon and then click **View Your Jobs**.
- e. In the Jobs panel, download your extracted meta, for example: **investigation-2020-May-19-08-30-20.json**.

- f. Go to back to the  (**Configure**) > **ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.

To download the events from the Investigate > Events view:

- a. In the main menu, go to **Investigate > Events** in a new tab.
 - b. In the Events view, enter a query for the ESA rule test.
 - c. Select the events to use and in the **Download** or **Download All** menu, select **Visible Meta as JSON** or **All Meta as JSON**, depending on the size of your selection.
 - d. In the main menu, go to **Dashboards** and in the toolbar click the  (**Jobs**) icon and then click **View Your Jobs**.
 - e. In the Jobs panel, download your extracted meta, for example: **Concentrator_ALL_EVENTS_ALL_META.json**.
 - f. Go to back to the  (**Configure**) > **ESA Rules** tab opened previously and copy the contents of the JSON file into the **Input Data** field in your ESA rule.
5. Click **Test Rule**. The **Output** field shows the output of your rule and you can determine if the results meet your requirements.

Test Rule Get immediate feedback on how the rule runs ^

ESA Service: ESA-ESA Correlation

Input Data:

```

    "Authentication"
  ],
  "ec.outcome": [
    "Failure"
  ],
  "reference.id": [
    "605004"
  ],
  "event.desc": [
    "Login denied"
  ]

```

Test Rule

Output: Test complete

- ✓ Rule successfully validated
- ✓ Provided input is valid
- ✓ Test ran successfully

Engine Stats

Engine Version	Events Offered	Offered Rate	Runtime Errors
8.4.0	9	0	-

Rule Stats

Deployed	Statements Fired	Alerts Fired	Events in Memory	Memory Usage	CPU %	Events Matched	Alerted Events	Runtime Errors	Debug Logs
✓	0	9	0	0	100	9	Details...	-	Details...

Note: If you are testing any Rule that has meta key defined as type 'short', the Test Rule will not generate alert for the event.

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Event Processing Language (EPL)

This topic describes Event Processing Language (EPL), a declarative language for dealing with high frequency time-based event data. ESA uses Event Processing Language (EPL), a declarative language for dealing with high frequency time-based event data. It is used for express filtering, aggregation, and joins over possibly sliding windows of multiple event streams. EPL also includes pattern semantics to express complex temporal causality among events. It can perform, but is not limited to, the following functions:

- Filter Event
- Alert Suppression
- Compute percentages or ratios
- Average, count, min and max for a given time window
- Correlate events arriving in multiple stream
- Correlate events that arrive out of order
- On-Off Windows
- Followed-by and Not Followed-by support
- Regex filter support

Databases require explicit querying to return meaningful data and are not suited to push data as it changes. The developer must implement the temporal and aggregation logic himself. By contrast, the EPL engine provides a higher abstraction and intelligence and can be thought of as a database turned upside-down. Instead of storing the data and running queries against stored data, EPL allows applications to store queries and continuously run the data through. Response from the EPL engine is real-time when conditions occur that match user defined queries.

For best practices on writing advanced EPL rules, see [ESA Rule Writing Best Practices](#).

Advanced ESA rules require correct character case, but in the Investigate Navigate view all characters are converted to lowercase. However, the meta may not be lowercase despite appearances in the Investigate Navigate view. To ensure you are using the correct case, you can make a strict pattern match for better performance.

Strict Pattern Match Example

```
@RSAAlert(oneInSeconds=0)
SELECT * FROM Event(
  (medium IN ( 1 ) AND
   filetype IN ( 'pdf' , 'windows_executable' , 'x86 pe' , 'windows
executable' ))
).win:time(5 Minutes)
MATCH_RECOGNIZE (
  MEASURES E1 as e1_data , E2 as e2_data
  PATTERN (E1 E2)
  DEFINE
    E1 as (E1.filetype IN ('pdf')),
    E2 as (E2.filetype IN ( 'pdf' , 'windows_executable' , 'x86 pe' ,
'windows executable' ))
```

```
);
```

Caution: Care should be taken to only add the case-insensitive *toLowerCase()* function on meta keys as needed. The *toLowerCase()* function can cause significant performance decreases. Consider checking the Investigate Events view or the Event Analysis view to see the real character case for meta fields and avoid unnecessary usage of the function.

For the purposes of online help, basic statements are used to illustrate how to set up ESA; however, for more information about writing EPL statements, the <http://www.espertech.com> site provides tutorials and examples.

Note: In NetWitness version 11.5, ESA Correlation supports Esper version 8.4.0. In version 11.4, ESA Correlation supports Esper version 8.2.0 and in version 11.3, ESA Correlation supports Esper version 7.1.0.

ESA Annotations

This topic describes annotations that NetWitness provides to use in advanced EPL rules.

For best practices on writing advanced EPL rules, see [ESA Rule Writing Best Practices](#).

@RSAContext Annotation (11.5 and later)

The @RSAContext annotation can be used in advanced rules to dynamically add or remove data from a Context Hub list after the rule fires. For example, you can create a rule that automatically adds an IP address to a blacklist and removes it from a whitelist.

You can update a single-column or a multi-column Context Hub list. The @RSAContext annotation also performs error handling when the Context Hub list cannot be reached.

Prerequisites

Before deploying a rule using the @RSAContext annotation, the list must exist in Context Hub. For information on creating a Context Hub list, see the *Context Hub Configuration Guide*. Go to the [NetWitness Documents - All Versions](#) to find all RSA NetWitness Platform 11.x documents.

IMPORTANT: If you rename a Context Hub list or recreate the Context Hub list with the same name, update the ESA rules that use that Context Hub list, and then redeploy the ESA rule deployments that contain those rules.

Single-Column Context Hub Lists

The @RSAContext annotation uses the following format for a single column Context Hub list:

```
@RSAContext (list="<single_column_list>", action=<DELETE_ENTRY or ADD_ENTRY>,
onError = <STOP_ALL_RULE_PROCESSING_AND_WAIT or IGNORE_ERROR_AND_CONTINUE>,
fields={"LIST=<meta_key>"})
```

The @RSA Context parameters for a **single-column** Context Hub List are described in the following table.

Parameter	Description
list	Where <single_column_list> is the name of the single-column Context Hub list (whitelist or blacklist).
action	<p>You can add an entry to or remove an entry from the list:</p> <ul style="list-style-type: none"> DELETE_ENTRY - Removes an entry from the list. ADD_ENTRY - Adds an entry to the list. <p>You can only have one action per @RSAContext entry.</p>

Parameter	Description
onError	<p>Identifies how to handle errors, for example, when the Context Hub list is not available or is full. You can choose one of the following options:</p> <ul style="list-style-type: none"> • <code>STOP_ALL_RULE_PROCESSING_AND_WAIT</code>: If any error occurs with the selected actions on the Context Hub list, stop all rule processing including data aggregation and retry indefinitely until the actions are successfully executed. Halting processing for one Context Hub list results in halting processing for all Context Hub lists. Select this option if it is important to update a Context Hub list before processing any more rules. For example, you may have a blacklist that must be updated before continuing rule processing. • <code>IGNORE_ERROR_AND_CONTINUE</code>: If any error occurs with the selected actions on the Context Hub list, ignore the error and the <code>@RSAContext</code> entry, and continue processing. If it is not a Context Hub error, assume that the error occurs before reaching Context Hub, and retry until the actions are successfully executed. Select this option if continuing rule processing is more important than updating the Context Hub list.
fields	Maps a meta key value to the column in the Context Hub list. Single-column lists always show <code>LIST=</code> in capital letters.

The first example deletes a source IP address from an IP whitelist. However, if the list is not available, the source IP address is not removed from the list and processing is continued.

```
@RSAContext (list="ip_whitelist", action=DELETE_ENTRY, onError = IGNORE_ERROR_AND_CONTINUE, fields={"LIST=ip_src"})
```

The second example adds a source IP address to an IP blacklist. However, if the list is not available, all rule processing stops including data aggregation, and the ESA Correlation service retries indefinitely until the entry is added to the specified Context Hub list.

```
@RSAContext (list="ip_blacklist", action=ADD_ENTRY, onError = STOP_ALL_RULE_PROCESSING_AND_WAIT, fields={"LIST=ip_src"})
```

Multi-Column Context Hub Lists

The `@RSAContext` annotation uses the following format for a multi-column Context Hub list:

```
@RSAContext (list="<multi_column_list>", action=<DELETE_ENTRY or ADD_ENTRY>,
onError = <STOP_ALL_RULE_PROCESSING_AND_WAIT or IGNORE_ERROR_AND_CONTINUE>,
fields={"<column1>=<meta_key1 or constant value1>",<column2>=<meta_key2 or constant value2>",<column n>=<meta_key n or constant value n>"},)
```

The `@RSAContext` parameters for a **multi-column** Context Hub List are described in the following table.

Parameter	Description
list	Where <multi_column_list> is the name of the multi-column Context Hub list.
action	<p>You can add an entry to or remove an entry from the list:</p> <ul style="list-style-type: none"> • DELETE_ENTRY - Removes an entry from the list. • ADD_ENTRY - Adds an entry to the list. <p>You can only have one action per @RSAContext entry.</p>
onError	<p>Identifies how to handle errors, for example, when the Context Hub list is not available or is full. You can choose one of the following options:</p> <ul style="list-style-type: none"> • STOP_ALL_RULE_PROCESSING_AND_WAIT: If any error occurs with the selected actions on the Context Hub list, stop all rule processing including data aggregation and retry indefinitely until the actions are successfully executed. Halting processing for one Context Hub list results in halting processing for all Context Hub lists. Select this option if it is important to update a Context Hub list before processing any more rules. For example, you may have a blacklist that must be updated before continuing rule processing. • IGNORE_ERROR_AND_CONTINUE: If any error occurs with the selected actions on the Context Hub list, ignore the error and the @RSAContext entry, and continue processing. If it is not a Context Hub error, assume that the error occurs before reaching Context Hub, and retry until the actions are successfully executed. Select this option if continuing rule processing is more important than updating the Context Hub list.
fields	Maps a meta key values or constant values to the columns in the Context Hub list. Multi-column lists show the names of each column in the fields parameter.*

The first example deletes a user and the associated source and destination IP addresses from a multi-column IP whitelist. However, if the list is not available, the specified fields are not removed from the list and processing is continued.

```
@RSAContext (list="MultColumn_whitelist", action=DELETE_ENTRY, onError =
IGNORE_ERROR_AND_CONTINUE, fields={"source=ip_src","destination=ip_
dst","user='smith'"})
```

The second example adds a user and the associated source and destination IP addresses to a multi-column IP blacklist. However, if the list is not available, all rule processing stops including data aggregation, and the ESA Correlation service retries indefinitely until the entry is added to the specified Context Hub list.

```
@RSAContext (list="MultColumn_blacklist", action=ADD_ENTRY, onError = STOP_
ALL_RULE_PROCESSING_AND_WAIT, fields={"source=ip_src","destination=ip_
dst","user='smith'"})
```

Automatic Context Hub List Updates

When using the `@RSAContext` annotation in your rules, string array meta keys (shown as `string[]` in the Meta Key References on the ESA Rules > Settings tab) are inserted and removed as separate lines in the Context Hub list. This enables lookups in the Investigate and Respond views.

Single Column Context Hub List Update Example

A single column Context Hub list has one column with the string array meta key `alias_host`. If the rule using that list fires and `alias_host` has three values (Google, Yahoo, and Dell), ESA Correlation adds three rows to the Context Hub list:

Google
Yahoo
Dell

Multi-Column Context Hub List Update Example

A multi-column Context Hub list has three columns. If the rule using that list fires and the meta keys have the following values:

- `username` (string array) with values Aimee and Chris
- `ip_dst` (string) with value 10.10.10.10
- `alias_host` (string array) with values Dell, Google, and Yahoo

ESA Correlation adds six rows to the Context Hub list:

Aimee, 10.10.10.10, Dell
Aimee, 10.10.10.10, Google
Aimee, 10.10.10.10, Yahoo
Chris, 10.10.10.10, Dell
Chris, 10.10.10.10, Google
Chris, 10.10.10.10, Yahoo

@RSAAlert Annotation

The `@RSAAlert` annotation is used to mark which EPL statements are linked to generating alert notifications. It is designed to work with the alert notification suppression feature in the Rule Builder user interface.

The `@RSAAlert` annotation can be useful when working with alert notifications, especially if you want to filter notifications, such as sending one notification for each user that triggers an alert.

For example, suppose you want to generate alert notifications for login failures. You could add the following statement:

```
@RSAAlert select * from event(msg_id="login_fail")
```

Event number	Message ID	username	src_IP	Time
1	login_fail	alice	1.2.3.4	10:00
2	login_fail	alice	1.2.3.4	10:01
3	login_fail	alice	6.7.8.9	10:01
4	login_fail	bob	1.2.3.4	10:01
5	login_fail	alice	1.2.3.4	10:03

For the above statement, five alert notifications are generated.

However, suppose you wanted to modify the statement to generate one alert for each separate username. You can use the *identifier* attribute. For example, the statement `@RSAAlert(identifier={"username"}) SELECT* FROM Event(msg_id="login_fail")` generates one notification for the first alert for "bob" and one for the first alert for "alice." Subsequent alerts for "bob" and "alice" are ignored.

You can further distinguish the users by adding details via the identifier variable. For example, you can distinguish by user and IP address using the following statement: `@RSAAlert(identifier={"username", "src_ip"}) SELECT* FROM Event(msg_id="login_fail")`. Then, you would see notifications generated by user name and IP address (one alert for "alice" at 1.2.3.4, another alert for "alice" at 6.7.8.9, and an alert for "bob" at 1.2.3.4).

To use identifiers with Alert Notification Suppression:

The `@RSAAlert` annotation is designed to work with the alert notification suppression feature in the Rule Builder user interface. To do this:

1. Create a rule in the Rule Builder user interface, and select the alert suppression feature when configuring notifications.

Output	Notification	Notification Server	Template
<input checked="" type="checkbox"/> SYSLOG	Local_SysLog	localhost-514	Default Syslog Template

Output Suppression of every minutes

2. Copy the code from the Rule Builder rule into a new advanced rule.
3. Configure the advanced rule to include identifiers (as described above) and save the advanced rule.
4. Delete the original rule builder rule.

@RSAPersist Annotation

The @RSAPersist annotation is used to mark a named window as an ESA managed window for persistence. By marking the named window as an ESA managed window, ESA periodically writes the contents of the window to disk and restores them back if the window is undeployed and redeployed. The systems take a snapshot just before the ESA rule deployment is undeployed and the window is removed. Conversely, it restores the window contents from the snapshot just after the deployment is redeployed. This ensures that the contents of the window are not lost if the deployment state is altered or if the ESA service goes down.

For example, consider a named window, DHCPTracker that holds a mapping from IP addresses to each assigned hostname. You can annotate the statement with the @RSAPersist annotation as:

```
@RSAPersist
  create window DHCPTracker.std:unique(ip_src) as (ip_src string, alias_
host string);
  insert into DHCPTracker select IP as ip_src, HostName as alias_host from
DHCPAssignment (ID=32);
```

Note: All windows definitions are not suitable for persistence. @RSAPersist annotation must be used with care. If the window has timed-records or if it depends on time based constraints it is very likely that the reverted snapshots will not restore it to the correct state. Also, any changes to the window definition will invalidate the snapshots and reset the window to a blank state. The system does not do any semantic analysis to determine if the changes to the window definition are conflicting or not. Note that other parts of a deployment (that is, other than the particular CREATE WINDOW call that defines the window) may change, without invalidating the snapshots.

Caution: (This caution applies only to NetWitness Platform versions 11.3.x, 11.4.x, and 11.5.0.0.) To avoid data being overwritten, if you have a rule with a named window, do not disable and re-enable it. Instead, undeploy and redeploy the ESA rule deployment that contains the rule.

@UsesEnrichment (10.6.1.1 and later)

The @UsesEnrichment can be used in advanced EPL rules to reference enrichments. In order to synchronize enrichments with ESA, all enrichment dependencies in EPL rules must be referenced with the @UsesEnrichment annotation.

The @UsesEnrichment annotation uses the following format:

```
@UsesEnrichment(name= '<enrichment_name>')
```

For example, the following EPL references a whitelist enrichment:

```
@UsesEnrichment(name = 'Whitelist')
@RSAAlert
SELECT * FROM Event(ip_src NOT IN (SELECT ip_address FROM Whitelist))
```

@Name

The @Name is the statement name defined in ESA advanced rules. It is used to dynamically generate statement names in ESA alerts. The statement name of only an alert triggering statement is displayed. This annotation has meta keys enclosed in curly brackets.

The @Name annotation uses the following format:

```
@Name("<static_part_of_statement_name> {meta_key1} {meta_key2}...")
```

For example, the following EPL references meta keys *ip_src* and *user_name* whose values will be dynamically generated.

```
@Name("Login Event to {ip_src} by {user_name}")
```

Note: You can specify any number of meta keys in the statement for dynamic statement name generation.
The length of individual meta key is limited to 64, after which the value is truncated and appended with "...".
The length of the dynamic generation of statement name is limited to 128, after which the value is truncated to 128 and appended with "...". All the remaining values post truncation will be treated as static values.

@Audit Annotation

Add the `@Audit('stream')` annotation to your ESA rules to print alerts to the ESA logs for troubleshooting. This is useful when debugging the Esper rules. The `@Audit('stream')` annotation provides debug information for the next statement in the rule. For example:

```
@Audit('stream')
```

```
@RSAAlert
```

```
SELECT * FROM Event((ip_src NOT IN ( '1.1.1.1' , '2.2.2.2' )));
```

In the above example, `@RSAAlert` is only necessary if the statement needs to send an alert to the Respond view.

In NetWitness Platform 11.5 and later, you can test rules in the rule builders. If you add the `@Audit('stream')` annotation to an advanced EPL rule, you can view the Debug Log in the test output. For more information, see [Validate an Advanced EPL Rule](#).

Example Advanced EPL Rules

Following are the examples of Advanced ESA rules. Each example has multiple ways of implementing the same use-case.

For best practices on writing advanced EPL rules, see [ESA Rule Writing Best Practices](#).

Example #1:

Create a user account and delete the same user account in 300s. User information is stored in user_src meta.

EPL #1:

Rule Name	CreateAndDelete Useraccount1
Rule Description	Create a user account followed by an action to delete the same user account in 300 seconds.
Rule Code	<pre>@RSAAlert (oneInSeconds=0) SELECT * FROM Event(ec_subject='User' AND ec_outcome='Success' AND user_src is NOT NULL AND ec_activity IN ('Create', 'Delete')) .win:time(300 seconds) match_recognize (partition by user_src measures C as c, D as d pattern (C D) define C as C.ec_activity='Create' , D as D.ec_activity='Delete');</pre>
Note	<ul style="list-style-type: none"> Filter events needed for pattern in given time frame. Filter conditions should be such that only required events are passed to match recognize function. In this case, they are create and delete user account Events. That is, Event(ec_subject='User' AND ec_outcome='Success' AND user_src is NOT NULL AND ec_activity IN ('Create', 'Delete')) Partition by creates buckets. In this case, Esper creates buckets per value of user_src. And hence value of user_src is common between both events. Define the pattern you want. Right now it is set to Create Followed by Delete. You can do multiple creates followed by delete (C+ D). Pattern is very similar to regular expression. Most efficient use case. The 'loose' pattern match of (C+ D) will result in decreased performance. Unless you need to include all C events within the generated alert, keep the strict pattern match of (C D). See the Esper documentation for more details.

EPL #2:

Rule Name	CreateAndDeleteUseraccount2
Rule Description	Create a user account followed by an action to delete the same user account in 300 seconds.
Rule Code	<pre>@RSAAlert (oneInSeconds=0) SELECT * from pattern[every (a= Event(ec_subject='User' AND ec_outcome='Success' AND user_dst is NOT NULL AND ec_activity IN ('Create')) -> (Event(ec_subject='User' AND ec_outcome='Success' AND user_dst is NOT NULL AND ec_activity IN ('Create') AND user_src = a.user_src)))where timer:within(300 Sec)];</pre>
Note	<ul style="list-style-type: none"> • Lets say same user is created twice and deleted once in that order. Then the above pattern will fire 2 alerts. • A thread is created for every User creation. • There is no way to control threads. It is important to have time bounds and preferably small intervals. • If you do not need every first event to start a new thread and match with the subsequent second event, then add suppression syntax of <code>@SuppressOverlappingMatches</code> after the pattern keyword. See the Esper documentation for more details.

Example #2:

Detect pattern where user created followed by login by same user and user is deleted in end. In case of windows logs user info is stored in either user_dst or user_src depending on event.

user_src(create) = user_dst(Login) = user_src(Delete)

EPL #3:

Rule Name	CreateUserLoginandDeleteUser
Rule Description	Detect a pattern where a user creates a User account followed by login by the same user followed by deletion of the User account.
Rule Code	<pre>@RSAAlert (oneInSeconds=0) SELECT * FROM Event(ec_subject='User' and ec_activity in ('Create','Logon','Delete') and ec_theme in ('UserGroup', 'Authentication') and ec_outcome='Success').win:time(300 seconds) match_recognize (measures C as c, L as l, D as d pattern (C L D) define C as C.ec_activity = 'Create',</pre>

Note	<pre>L as L.ec_activity = 'Logon' AND L.user_dst = C.user_src, D as D.ec_activity = 'Delete' AND D.user_src = C.user_src);</pre>
	<ul style="list-style-type: none"> • Since user_src/user_dst is not common across all events we can't use partition. It will be 1 single bucket running 1 pattern at a time. For example, for user 1 and 2 if the stream of events are C1C2L1D1, C1L1C2D1, there will be no alert because C1 thread got reset by C2. Alert will be fired only if C1L1D1 are in order and no other event either from same user or other user falls in between. • Another solution would be to use Named Window and merge user_dst and user_src into single column and then run match recognize. (EPL #3). • Pattern can also be used. You might get more alerts than expected. (EPL #4).

EPL #4: Using NamedWindows and match recognize

Rule Name	CreateUserLoginandDeleteUser
Rule Description	Detect a pattern where a user creates a User account followed by login by the same user followed by deletion of the User account.
Rule Code	<pre>@Name('NormalizedWindow') create window FilteredEvents.win:time (300 sec) (user String, eactivity string, sessionid Long); @Name('UsersrcEvents') Insert into FilteredEvents select user_src as user, ec_activity as eactivity, sessionid from Event(ec_subject='User' and ec_activity in ('Create','Delete') and ec_theme in ('UserGroup', 'Authentication') and ec_ outcome='Success' and user_src is not null); @Name('UsrdstEvents') Insert into FilteredEvents select user_dst as user, ec_activity as eactivity, sessionid from Event(ec_subject='User' and ec_ activity in (Logon') and ec_theme in ('UserGroup', 'Authentication') and ec_outcome='Success' and user_dst is not null); @Name('Pattern') @RSAAlert(oneInSeconds=0, identifiers={"user"}) select * from FilteredEvents match_recognize (partition by user measures C as c, L as l, D as d pattern (C L+D) define C as C.eactivity= 'Create', L as L.eactivity= 'Logon', D as D.eactivity='Delete');</pre>

EPL #5: Using Every @RSAAlert(identifiers={"user_src"})

Rule Name	CreateUserLoginandDeleteUser
-----------	------------------------------

Rule Description	Detect a pattern where a user creates a User account followed by login by the same user followed by deletion of the User account.
Rule Code	<pre>SELECT a.time as time,a.ip_src as ip_src,a.user_dst as user_dst,a.ip_dst as ip_dst,a.alias_host as alias_host from pattern [every (a=Event (ec_subject='User' and ec_activity='Create' and ec_theme='UserGroup' and ec_outcome='Success') -> (Event(ec_subject='User' and ec_activity='Logon' and ec_theme='Authentication' and user_src=a.user_dst) -> b=Event(ec_subject='User' and ec_activity='Delete' and ec_theme='UserGroup' and user_dst=a.user_dst))) where timer:within(300 sec)];</pre>

Example #3:

Excessive login failures from same sourceIP.

EPL #6: @RSAAlert(identifiers={"ip_src"})

Rule Name	ExcessLoginFailure																																
Rule Description	The same user tried logging in from the same Source IP and faced login failures.																																
Rule Code	<pre>@RSAAlert(oneInSeconds=0) SELECT * FROM Event (ip_src IS NOT NULL AND ec_activity ='Logon' AND ec_outcome = 'Failure').win:time_batch(300 seconds) GROUP BY ip_src HAVING COUNT(*) = 10;</pre> <ul style="list-style-type: none"> • Uses time_batch: Looks at events in batches(tumbling window). Every event matching the filter criteria will be kept for the specified time window. • “GROUP BY” clause aggregates events within the data window by ip_src and HAVING clause instructs a count of 10 events with the same ip_src must occur within the time window. • One of the issues with tumbling windows is that events occurring towards the end of the batch might not lead to an alert. 																																
Note	<p>In the below sequence of events at t=301 even though 10 login failures occurred for the same login in the last 300 secs, there will be no alert because the batch of events was dropped at t=300.</p> <table border="1"> <thead> <tr> <th>Time t</th> <th>Login Failures for Specific Users</th> <th>Alert</th> <th>Time Batch</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>1</td> </tr> <tr> <td>295</td> <td>6</td> <td>0</td> <td>1</td> </tr> <tr> <td>299</td> <td>3</td> <td>0</td> <td>1</td> </tr> <tr> <td>301</td> <td>1</td> <td>0</td> <td>2</td> </tr> <tr> <td>420</td> <td>6</td> <td>0</td> <td>2</td> </tr> <tr> <td>550</td> <td>3</td> <td>0</td> <td>2</td> </tr> <tr> <td>600</td> <td>0</td> <td>0</td> <td>3</td> </tr> </tbody> </table>	Time t	Login Failures for Specific Users	Alert	Time Batch	0	0	0	1	295	6	0	1	299	3	0	1	301	1	0	2	420	6	0	2	550	3	0	2	600	0	0	3
Time t	Login Failures for Specific Users	Alert	Time Batch																														
0	0	0	1																														
295	6	0	1																														
299	3	0	1																														
301	1	0	2																														
420	6	0	2																														
550	3	0	2																														
600	0	0	3																														

720	6	0	3
850	3	0	3
900	1	1	3 ends and 4 begins

- Above problem can be resolved using win:time windows (EPL#7) instead of win:time_length_batch windows.
- Outer group by is to control events when time elapses. Say you have 9 events at end of 60 secs, Esper engine will push those 9 events to listener. Group by and count will restrict it since count is not equal to 10.
- Time and count can be modified as needed.

EPL #7: @RSAAAlert(identifiers={"ip_src"})

Rule Name	ExcessLoginFailure
Rule Description	The same user tried logging in from the same Source IP and faced login failures.
Rule Code	<pre>@RSAAAlert(oneInSeconds=0) SELECT * FROM Event (ip_src IS NOT NULL AND ec_activity ='Logon' AND ec_outcome = 'Failure').win:time(300 seconds) GROUP BY ip_src HAVING COUNT(*) = 10;</pre>
Note	<ul style="list-style-type: none"> • This is a sliding window and hence after an alert is fired for a set of events they can be used for another alert as well until time has passed. • If 10 events were involved in causing the alert only the last event will appear. • Events are not removed from the time window. You could use output rate limiting. See the Esper documentation for more details.

Example #4:

Multiple failed logins from multiple different users from same source to same destination, a single user from multiple different sources to same destination.

EPL #8: using time_batch

Rule Name	MultiplefailedLogins
Rule Description	<p>There are multiple failed logins for the following cases:</p> <ul style="list-style-type: none"> - From multiple users from same source to same destination. - Single user from multiple sources to the same destination.
Rule Code	<pre>@RSAAAlert(oneInSeconds=0) SELECT * FROM Event (ec_activity='Logon' AND ec_outcome='Failure'</pre>

Note	<pre> AND ip_src IS NOT NULL AND ip_dst IS NOT NULL AND user_dst IS NOT NULL) .win:time_batch(300 seconds) group by ip_src,ip_dst having count(distinct user_dst) >= 5; </pre>
	<ul style="list-style-type: none"> • ip.dst and ip.src are common across all events. • user_dst is unique for all events. • Alert is fired when there are at least 5 different users try to login from same ip.src and ip.dst combination.

Example #5:

No Log traffic from a device in a given timeframe.

EPL #9: using timer:interval

Rule Name	NoLogTraffic
Rule Description	There is no log traffic observed from a device in a given time frame.
Rule Code	<pre> SELECT * FROM pattern [every a = Event(device_ip IN ('10.0.0.0', '10.0.0.1') AND medium = 32) -> (timer:interval (3600 seconds) AND NOT Event(device_ip = a.device_ip AND device_ type = a.device_type AND medium = 32))]; </pre>
Note	<ul style="list-style-type: none"> • Rule only detects sudden loss of traffic. It won't alert if there is no traffic to begin with. You need at least 1 event for rule to alert. • List of device ip address or device hostnames as input. Only these systems will be tracked. • Time input is required. Alert is fired when time interval between events exceeds input time.

Example #6:

Multiple Failed Logins NOT followed by a Lockout event by the same user.

EPL #10: using timer and Lockout

Rule Name	FailedloginswoLockout
Rule	There are multiple failed logins that are not followed by Lockout event by the same

Description	user.
Rule Code	<pre>SELECT * FROM pattern [every-distinct(a.user_dst, a.device_ip, 1 msec) (a= Event(ec_activity='Logon' and ec_ outcome='Failure' and user_dst IS NOT NULL) -> [2](Event (device_ip =a.device_ip and ec_activity='Logon' and ec_ outcome='Failure' and user_dst=a.user_dst) AND NOT Event((ec_activity='Logon' and ec_ outcome='Success' and device_ip = a.device_ip and user_ dst=a.user_dst) or (ec_activity='Lockout' and device_ip = a.device_ip and user_dst=a.user_dst)))) where timer:within(60 seconds) -> (timer:interval(30 seconds) and not Event(device_ip=a.device_ip and user_ dst=a.user_dst and ec_activity='Lockout'))];</pre>
Note	<ul style="list-style-type: none"> • Above query detects the absence of a Lockout Event after the occurrence of 2 failed logins from same user. • The occurrence of the multiple failed logins are timed and are assumed to occur within a certain period of time. Also, in-practice the Lockout event is assumed to occur within a short time after the occurrence of the last failed login event because the threshold value of Failed logins per user is set in a given domain. • In current query, every distinct will suppress new thread for combination of user and device for 1 millisecc. • Time allowed for 3 failed logins is 60 secs since 1st failed attempt. Wait period for lockout event to occur is 30 secs

Example #7:

Custom functions to perform LIKE and REGEX operations for ARRAY elements.

EPL #11: @RSAAlert(oneInSeconds=0)

Rule Name	MatchLikeRegex
Rule Description	There are custom functions to perform LIKE and REGEX comparisons of array meta keys.
Rule Code	<pre>SELECT * FROM pattern[e1=Event(matchLike(alias_host, "10.0.0.%")) AND e2=Event(matchRegex(alias_host, "10\.0\.0\.1[0-9][0-9]")) where timer:within(5 Minutes)];</pre>

Note:

1. “.” in meta keys should be replaced with (“_”).
2. All patterns should be time bound.
3. Use appropriate tags in front of statements, for example:

```
@RSAPersist:  
@RSAAAlert:
```

For additional details you can refer to:

- EPL Documentation: <http://www.espertech.com/esper/esper-documentation/>
- EPL Online Tool: <http://esper-epl-tryout.appspot.com/epltryout/mainform.html>

Configure an In-Memory Table Using an EPL Query


Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules. Recurring In-Memory Tables are no longer supported; use Content Hub Lists as enrichment sources. For more information, see [Configure a Context Hub List as an Enrichment Source](#).

When you use an In-Memory Table configuration in expert mode, you can create an enrichment source or named window based on an Esper query. This allows you to have more control over the content and create more dynamic content. When you do this, an EPL query constructs the named window to capture interesting states from the event stream.

Workflow

The following shows the workflow for creating a query using a named window:

1. The event is sent to the Esper Engine.
2. An EPL query is generated.
3. An alert is triggered.
4. The query checks to see if there is a connection between the event and the Named Window.
5. If there is a connection, the query that populates the Named Window is run and populated.
6. The content from the Named Window is added to the alert content and sent or displayed (depending on your settings).





Prerequisites

- The meta used in the EPL statement must exist in the data.
- You must create well-formed EPL statements.

Procedure


Note: It is preferable to use Context Hub List enrichment sources instead of In-Memory Table enrichment sources for rules.

1. Go to  (**Configure**) > **ESA Rules**.
The Configure view is displayed with the Rules tab open.
2. Click the **Settings** tab.
3. In the options panel, select **Enrichment Sources**.
4. In the **Enrichment Sources** section, click  > **In-Memory Table**.

In-Memory Table

Upload Type: Adhoc Recurring

Enable

User-Defined Table Name* 

Description

Import Data

Expert Mode

Table Columns **+ -**

	Name	Type
<input type="checkbox"/>		

Key

Max Rows

Persist Stored File Format: Object JSON

[For information on how to define and use an In-Memory Table, see the documentation](#)

5. Select **Adhoc**.
By default, Enable is selected. When you add the in-memory table to a rule, alerts will be enriched with data from it.
6. In the **User-Defined Table Name** field, type a descriptive name to describe the in-memory table.
7. If you want to explain what the enrichment adds to an alert, enter information in the **Description** field.
This description displays when you view the list of enrichments from the Enrichment Sources view, so it's a good idea to enter a thorough description as a best practice. Doing this allows other users to understand the content of the enrichment without opening it to examine its contents.
8. Select **Expert Mode** to define an advanced in-memory table configuration by writing an EPL query. The Table Columns are replaced by a **Query** field.
9. Select **Persist** to preserve the in-memory table on disk when the ESA service stops and to re-populate the table when the service restarts.
10. Enter the EPL query in the **Query** field. The query should be well-formed, and it's a good idea to test it before entering it in the field.
11. Click **Save**.

Example

For example, you want to know when an IPS or IDS is giving five or more inbound events with an event identified with malicious code. Additionally, you would like to know when the source IP of those events has been identified as suspicious by other sources. This information helps to more quickly triage the event and determine whether the alert is a true positive.

Step 1: Create the Enrichment

In this example, this enrichment is a watchlist of IPs that have been identified as suspicious by third party sources or by internal staff. The meta of `threat_desc` equal to `'suspicious ip'` is generated when a match to a feed occurs. This meta can be matched and output based on a log, packet, or endpoint event.

The enrichment should look like the following:

In-Memory Table ✕

Upload Type: Adhoc Recurring

Enable

User-Defined Table Name*

Description

Import Data

Expert Mode

Query*

```
create window IpWatchlist .std:unique(ip_src) as (ip_src string, threat_source string,
threat_category string);

insert into IpWatchlist
select ip_src, threat_source, threat_category from Event
where threat_desc = 'suspicious ip';
```

[For information on how to define and use an In-Memory Table, see the documentation](#)

Parameters	Description
Upload Type	Adhoc
IP_Watchlist	IP_Watchlist
Description	Dynamically populated whitelist based on a feed of IPs that are considered suspicious.
Expert Mode	Selected
Query	<pre>create window IpWatchlist .std:unique(ip_src) as (ip_src string, threat_source string, threat_ category string); insert into IpWatchlist select ip_src, threat_source, threat_category from Event where threat_desc = 'suspicious ip';</pre>

Step 2: Create Your Rule

First, you need to create your ESA Correlation rule. This example rule looks for inbound IPS or IDS log events with the `event_cat_name` beginning with `Attacks.Malicious Code`. If five or more events for the same `ip_src` occur within 60 minutes, then an alert will be triggered. If an `ip_src` from the Enrichment equals the `ip_src` from the alert, then that alert will be enriched with additional meta. In this case, the analyst would see the values for `threat_source` and `threat_category` in the raw alert. `Threat_category` would indicate the type of malware and `threat_source` would indicate the entity that has reported the ip as suspicious. The analyst could use this information to do additional research or escalate to the next tier for creation of a possible incident.

Rule Statement

Build a Statement ? X

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

if all conditions are met + -

Key	Operator	Value	Ignore Case?	Array?
<input type="checkbox"/> event.medium	is	32	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> event.device_class	is	IPS, IDS	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> event.direction	is	inbound	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> event.event_cat_name	begins with	Attacks.Malicious Code	<input type="checkbox"/>	<input type="checkbox"/>

Rule Logic with Enrichment Added

Rules
Services
Settings
IDS or IPS Events with Mali... ✕

Rule Builder

Build a rule using drag-and-drop and auto-complete tools.

Rule Name *

Description

Trial Rule

Alert

Severity *

Conditions * [Investigation](#)

Statement	Occurs	Connector	Correlation Type	Meta	Meta
<input type="checkbox"/> IDS or IPS Events with Malicious Code	5				

Group By

Occurs Within minutes

Notifications [Global Notifications](#)

Output	Notification	Notification Server	Template
No parameters to edit.			

Output Suppression of every minutes

Enrichments [Settings](#)

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
<input checked="" type="checkbox"/> In-Memory Table	Ip_Watchlist	ip_src	ip_src

ESA Alert References

In Event Stream Analysis (ESA), you configure and deploy ESA rules to get alerted about potential network threats.

These topics explain the user interface for ESA Correlation rules.

- [Rules Tab](#)
- [Rules Tab Options Panel](#)
- [Rule Library Panel](#)
- [Rule Builder Tab](#)
- [Build a Statement Dialog](#)
- [Advanced EPL Rule Tab](#)
- [Rule Syntax Dialog](#)
- [Deployment Panel](#)
- [Deploy ESA Services Dialog](#)
- [Deploy ESA Rules Dialog](#)
- [Updates to the Deployment Dialog](#)
- [Services Tab](#)
- [Settings Tab](#)

Rules Tab

The Rules tab enables you to configure ESA rules and deployments.


What do you want to do?

Role	I want to ...	Show me how
Content Expert	View types of rules.	ESA Rule Types
Content Expert	Deploy Trial Rules.	Working with Trial Rules
Content Expert	Create a rule.	Add Rules to the Rule Library
Content Expert	Deploy a rule.	Deploy Rules to Run on ESA


Related Topics

- [Getting Started with ESA](#)

Quick Look

The Rules tab is displayed when you go to  **(Configure) > ESA Rules.**

The following figure shows the Rules tab.



The screenshot displays the RSA NetWitness Platform interface for the Rules tab. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation menu on the left shows 'Rules', 'Services', and 'Settings'. The 'Rule Library' section is active, showing a table of rules. The table has the following columns: Rule Name, Description, Trial Rule, Type, and Actions. The first rule is selected:

Rule Name	Description	Trial Rule	Type	Actions
<input type="checkbox"/> Endpoint Risk Scoring Rule Bundle	This bundle contains many rules which are used for risk scoring of E...	No	Endpoint Rule B...	
<input checked="" type="checkbox"/> SAMPLE - Blacklist - From inside countries that are not the US, N...	Monitors for non-SMTP traffic on TCP destination port 25 contai...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - Non SMTP Traffic on TCP Port 25 Containing Executable	Monitors for non-SMTP traffic on TCP destination port 25 containing ...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - P2P Software as Detected by an Intrusion Detection Device	P2P software as detected by an intrusion detection device (IDS), intru...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - User Added to Admin Group Same User su sudo	Alert when user is upgraded to one of admin groups (custom list of g...	Yes	Rule Builder	
<input type="checkbox"/> SAMPLE - Whitelist - From outside of Germany, P2P Software as Dete...	Whitelist Germany from P2P software as detected by an intrusion de...	Yes	Rule Builder	

The interface also includes a search filter, a table of actions, and a footer with 'Page 1 of 1' and 'Page Size 100'.

The Rules tab is divided into three sections:

- [Rules Tab Options Panel](#)
- [Rule Library Panel](#)
- [Deployment Panel](#)

Rules Tab Options Panel

In the **Rules** tab options panel to the left, you can view ESA rules in the Rule Library and create ESA rule deployments.

What do you want to do?

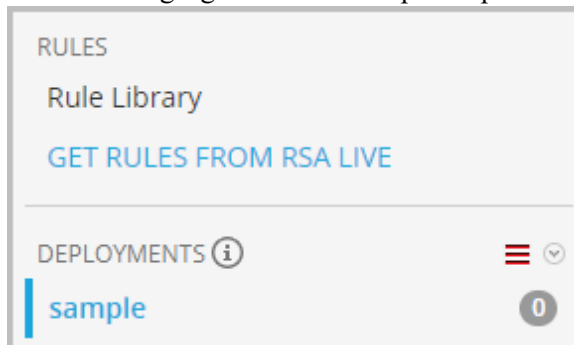
Role	I want to ...	Show me how
Content Expert	View an ESA rule.	Add Rules to the Rule Library
Content Expert	Create an ESA rule deployment.	ESA Rule Deployment Steps

Related Topics

- [Working with Rules](#)

Quick Look

The following figure shows the options panel in the **Rules** tab.






There are two sections in the options panel: Rules and Deployments.

Rules Section

The Rules section contains two options. **Rule Library** is selected by default, and when it's selected, the Rule Library view is displayed within the tab. **Get Rules From RSA Live** navigates to the Live Search view, where you can search for rules.

Deployments Section

The Deployments section lists ESA rule deployments and indicates whether there are updates to the deployments. From this section, deployments can be added, deleted, edited, and refreshed. Selecting a deployment from the list displays the Deployment panel within the tab. The following table describes the features of this section.

Feature	Description
	Displays a drop-down menu from which you can choose to add, edit, or delete an ESA rule deployment. You can also refresh the list of deployments to see if there are any new updates to the list.
	Indicates whether there are any updates to the deployment.
	Indicates the number of rules in the deployment.

Rule Library Panel

The Rule Library panel allows you to manage rules.

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Add an ESA rule.	Add a Rule Builder Rule
Content Expert	Edit, duplicate, or delete an ESA rule.	Edit, Duplicate or Delete a Rule
Content Expert	Import or export ESA rules.	Import or Export Rules
Content Expert	Filter the ESA rules list.	Filter or Search for Rules


Related Topics

- [Add an Advanced EPL Rule](#)

Quick Look

To access this view, go to  **(Configure) > ESA Rules**. The Rules tab is displayed and the Rule Library panel is on the right.

The following figure shows the Rule Library panel.



Rule Library
All Event Stream Analysis (ESA) rules created or downloaded. Add rules to a deployment to activate them.

Filter

<input type="checkbox"/>	Rule Name	Description	Trial Rule	Type	Actions
<input type="checkbox"/>	SAMPLE - Blacklist - From inside countries that are not the U...	Monitors for non-SMTP traffic on TCP desti...	Yes	Rule Builder	
<input type="checkbox"/>	SAMPLE - Non SMTP Traffic on TCP Port 25 Containing Exec...	Monitors for non-SMTP traffic on TCP desti...	Yes	Rule Builder	
<input type="checkbox"/>	ESA - Recon Enrichment	test	Yes	Rule Builder	
<input type="checkbox"/>	ESA: ip.src not null - Custom Enrichment	, wolverine, adamantium_claws ...	No	Rule Builder	
<input type="checkbox"/>	ESA - GeoIP	This is not a trial rule	No	Rule Builder	
<input type="checkbox"/>	ESA Events for every source ip	Description from ESA Rule goes here.	No	Rule Builder	
<input type="checkbox"/>	ESA - IP Enrichment Data	This is a test	No	Rule Builder	
<input type="checkbox"/>	ESA - In memory enrichment	Enrichment data from csv	No	Rule Builder	
<input type="checkbox"/>	UserName Enrichment	Enrichment data from csv	No	Rule Builder	

Page 1 of 1 | Page Size 100 | Displaying 1 - 12 of 12 rules

The Rule Library panel includes the following components:

- Rule Library toolbar
- Rule Library list

Rule Library Toolbar

The Rule Library toolbar allows you to add, delete, edit, duplicate, filter, export, and import ESA rules. The following figure shows the icons for these actions.



Rule Library List

The following figure shows the Rule Library list.

<input type="checkbox"/>	Rule Name	Description	Trial Rule	Type	Actions
<input type="checkbox"/>	SAMPLE - Blacklist - From inside countries that are not the U...	Monitors for non-SMTP traffic on TCP desti...	Yes	Rule Builder	
<input type="checkbox"/>	SAMPLE - Non SMTP Traffic on TCP Port 25 Containing Exec...	Monitors for non-SMTP traffic on TCP desti...	Yes	Rule Builder	
<input type="checkbox"/>	ESA - Recon Enrichment	test	Yes	Rule Builder	
<input type="checkbox"/>	ESA: ip.src not null - Custom Enrichment	██████████, wolverine, adamantium_claws ...	No	Rule Builder	
<input type="checkbox"/>	ESA - GeoIP	This is not a trial rule	No	Rule Builder	
<input type="checkbox"/>	ESA Events for every source ip	Description from ESA Rule goes here.	No	Rule Builder	
<input type="checkbox"/>	ESA - IP Enrichment Data	This is a test	No	Rule Builder	
<input type="checkbox"/>	ESA - In memory enrichment	Enrichment data from csv	No	Rule Builder	
<input type="checkbox"/>	UserName Enrichment	Enrichment data from csv	No	Rule Builder	


Page 1 of 1 | Page Size 100 | Displaying 1 - 12 of 12 rules

The Rule Library list shows all of the ESA rules. The following table lists the columns in the Rule Library list and their description.

Column	Description
Rule Name	Purpose of the ESA rule.
Description	Summary of what the ESA rule detects.
Trial Rule	Deployment mode to see if the rule runs efficiently.
Type	The type of rule. For more information, see ESA Rule Types .
Actions ()	Menu to delete, edit, duplicate, or export the selected rule.
Severity	Threat level of alert triggered by the rule.
Email	Indicates whether an alert notification for the rule is sent by email. This column is not visible by default.

Column	Description
SNMP	Indicates whether an alert notification for the rule is sent using SNMP. This column is not visible by default. (ESA SNMP notifications are not supported in NetWitness version 11.3 and later.)
Syslog	Indicates whether an alert notification for the rule is sent using Syslog. This column is not visible by default.
Script	Indicates whether an alert notification for the rule executes a script. This column is not visible by default.
Last Modified	The date and time when the ESA rule was last modified. This column is not visible by default.

To display columns which aren't visible by default, hover over the title of a column and click the **v** on the right. This opens a drop-down menu in which you can sort the contents of the column or choose which columns you want to see in the Rule Library list.



Rule Builder Tab

The Rule Builder tab enables you to define a Rule Builder rule.

What do you want to do?



Role	I want to ...	Show me how
Content Expert	Define a Rule Builder rule.	Add a Rule Builder Rule Step 1. Name and Describe the Rule
Content Expert	Define rule criteria.	Step 2. Build a Rule Statement
Content Expert	Add conditions to the rule.	Step 3. Add Conditions to a Rule Statement
Content Expert	Test the ESA rule logic.	Validate an ESA Rule

Related Topics

- [Add an Advanced EPL Rule](#)

Quick Look

To access the Rule Builder tab:

1. Go to  (**Configure**) > **ESA Rules**.
The Rules tab opens by default.
2. In the **Rule Library** toolbar, select  > **Rule Builder**.
The Rule Builder tab is displayed.

The following figure shows the Rule Builder tab.

Rule Builder
Build a rule using drag-and-drop and auto-complete tools.

Rule Name *

Description

Trial Rule

Memory Threshold MB

Alert

Severity *

Conditions * [Investigation](#)

Statement	Occurs	Connector	Correlation Type	Meta	Meta
Please add at least one statement.					

Group By

Occurs Within minutes

Notifications [Global Notifications](#)

Output	Notification	Notification Server	Template
No parameters to edit.			

Output Suppression of every minutes

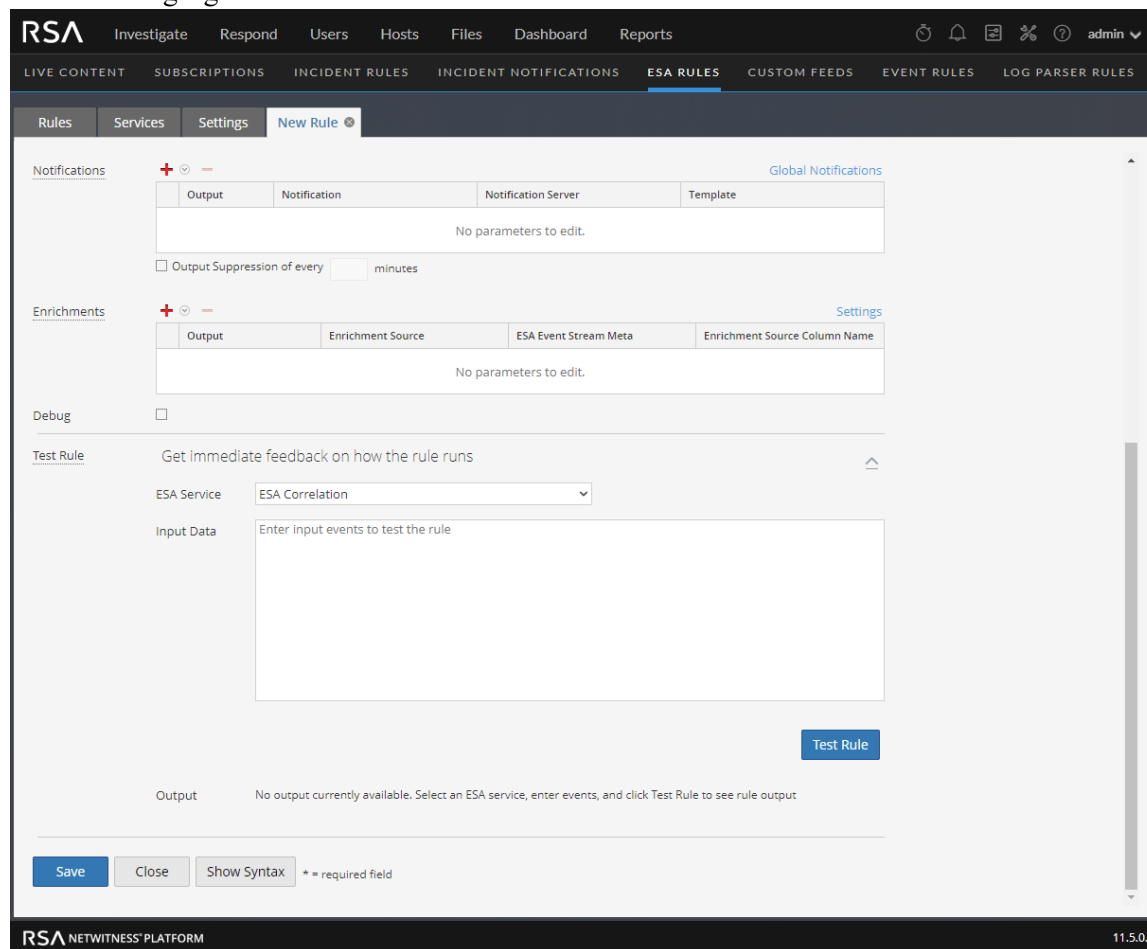
Enrichments [Settings](#)

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
No parameters to edit.			


Debug

RSA NETWITNESS PLATFORM 11.5.0.0

The following figure shows the Rule Builder tab scrolled down with the Test Rule section in view.



The following table lists the parameters in the Rule Builder tab.

Field	Description
Rule Name	Purpose of the ESA rule.
Description	Summary of what the ESA rule detects.
Trial Rule	Deployment mode to see if the rule runs efficiently.
Memory Threshold	(This option applies to version 11.5 and later.) The maximum memory usage allowed for this rule in MB. Add Memory Thresholds to ESA rules that use memory. For example, if a rule contains windows or pattern matching, configure a memory threshold for that rule. If the configured memory threshold is exceeded, it gets disabled individually and an error is displayed for that rule on the  (Configure) > ESA Rules > Services tab. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.

Field	Description
Alert	(This option applies to version 11.3 and later.) When selected, the alert is sent to Respond. If the checkbox is cleared, an alert will not be sent to Respond. To turn alerts on or off for ALL rules, see the <i>ESA Configuration Guide</i> .
Severity	Threat level of alert triggered by the rule.

The Rule Builder includes the following components:

- [Conditions Section](#)
- [Notifications Section](#)
- [Enrichments Section](#)
- [Debug Option](#)
- [Test Rule Section](#)

Conditions Section

In the Conditions section of the Rule Builder tab, you define what the rule detects.

The following figure shows the Conditions section.

The following table lists the parameters of the Conditions section.

Parameter	Description
+	Add a statement.
-	Remove selected statement.
	Edit selected statement.
Statement	Logical group of conditions for one operation.
Occurs	Alert frequency if the condition is met. This specifies that there must be at least that many events that satisfy the criteria in order to trigger an alert. The time window in minutes binds the Occurs count.

Parameter	Description
Connector	Options to specify relationship among the statements: <ul style="list-style-type: none"> • followed by • not followed by • AND • OR <p>The Connector joins two statements with AND, OR, followed by, or not followed by. When followed by is used, it specifies that there is a sequencing of those events. AND and OR build one large criteria. The followed by creates distinct criteria that occurs in sequence.</p>
Correlation Type	Correlation Type applies only to followed by and not followed by . If you choose a correlation type of SAME, select one meta to correlate on, and if you choose a correlation type of JOIN, select two meta to correlate on. You may want to use JOIN if you are trying to correlate on meta from two different data sources. For example, say you want to correlate an AV alert with an IDS alert.
Meta	Enter the meta condition if choosing a correlation type of SAME or JOIN (as described above).
Meta	Enter the second meta condition if choosing a correlation type of JOIN (as described above). For example, The destination IP address from the AV alert and source IP address for the workstation from the IDS alert are joined to allow you to view the same entities across different sources.
occurs within minutes	Time window within which the conditions must occur.
Event Sequence	Choose whether the pattern must follow a <i>strict</i> match or a <i>loose</i> match. If you specify a strict match, this means that the pattern must occur in the <i>exact</i> sequence you specified with no additional events occurring in between. <p>For example, if the sequence specifies five failed logins (F) followed by a successful login (S), this pattern will only match if the user executes the following sequence: F,F,F,F,F,S. If you specify a loose match, this means that other events may occur within the sequence, but the rule will still trigger if all of the specified events also occur. For example, five failed login attempts (F), followed by any number of intervening successful login attempts (S), followed by a successful login attempt might create the following pattern: F,S,F,S,F,S,F,S,F,S which would trigger the rule despite the intervening successful logins.</p>
Group By	Select the meta key by which to group results from the dropdown list. <p>For example, suppose that there are three users; Joe, Jane, and John and you use the Group By meta, user_dst (user_dst is the meta field for the user destination account). The result will show events grouped under the user destination accounts, Joe, Jane, and John.</p> <p>You can also group by multiple keys. For example, you might want to group by user and machine to see if a user logged in from the same machine attempts to log into an account multiple times. To do this, you might group by user_dst and ip_src.</p>

Notifications Section

In the Notifications section, you can choose how to be notified when ESA generates an alert for the rule. For more information on the alert notifications, see [Add Notification Method to a Rule](#).

The following figure shows the Notifications section.

Output	Notification	Notification Server	Template
<input checked="" type="checkbox"/> SYSLOG	Local_SysLog	localhost-514	Default Syslog Template

Output Suppression of every minutes

Parameter	Description
	To add an alert notification type.
	To delete the selected alert notification.
Output	Alert notification type. Options are: <ul style="list-style-type: none"> Email SNMP (This option is not supported in NetWitness version 11.3 and later.) Syslog Script
Notification	Name of previously configured output, such as an email distribution list.
Notification Server	Name of server that sends the output.
Template	Name of template for the alert notification.
Output Suppression of every	Option to specify alert frequency.
Minutes	Alert frequency in minutes.



Enrichments Section

In the Enrichments section, you can add a data enrichment source to a rule.

For more information on the enrichments, see [Add an Enrichment to a Rule](#).

The following figure shows the Enrichments section.

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
<input type="checkbox"/> In-Memory Table	Select Enrichment Source	Enter Meta	Enter Column Name
<input checked="" type="checkbox"/> GeoIP	Select Enrichment Source	Enter Meta	ipv4

Parameter	Description
	To add an enrichment.
	To delete the selected enrichment.
Output	Enrichment source type. Options are: <ul style="list-style-type: none"> In-Memory Table (Ad hoc only - Recurring In-Memory Tables are no longer supported in version 11.3 and later.) GeoIP
Enrichment Source	Name of previously configured enrichment source, such as a .CSV filename for an In-Memory Table.
ESA Event Stream Meta	ESA meta key whose value will be used as one operand of join condition.
Enrichment Source Column Name	Enrichment source column name whose value will be used as the other operand of the join condition. For an in-memory table, If you configured a key when creating a .CSV-based enrichment, this column automatically populates with the selected key. However, you can change it if you like. For a GeoIP enrichment source, ipv4 is automatically selected.

Debug Option

Select the Debug option to print alerts to the ESA logs for troubleshooting. This adds an @Audit ('stream') annotation to the rule. This is useful when debugging the Esper rules.

Test Rule Section

Note: The Test Rule section is available in NetWitness Platform 11.5 and later.

In the Test Rule section, you can validate your ESA rule to determine if the rule logic is working as expected before deploying the rule.

Test Rule Get immediate feedback on how the rule runs ^

ESA Service: ESA- ESA Correlation

Input Data:

```

],
"ip.dst": [
  "10.100.10.1"
],
"direction": [
  "inbound"
],
"service.name": [
  "telnet"
],
],
```

Test Rule

Output: Test complete

- ✔ Rule successfully validated
- ✔ Provided input is valid
- ✔ Test ran successfully

Engine Stats

Engine Version	Events Offered	Offered Rate	Runtime Errors
8.4.0	5870	0	-

Rule Stats

Deployed	Statement Fired	Alerts Fired	Events in Memory	Memory Usage	CPU %	Events Matched	Alerted Events	Runtime Errors	Debug Logs
✔	0	5870	0	0	100	5870	Details...	-	Details...

Field	Description
ESA Service	Select the ESA Correlation service to process the rule.
Input Data	Enter the input events to test the rule. Download the events from the Investigate view in JSON format, copy the events, and paste them in this field.
Output Data	After you select an ESA Correlation service, input data, and click the Test Rule button, you can view the output of the rule here and verify that the rule is working according to your requirements. You can view the alerts in the output, but this test does not send any alert notifications.

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service

Field	Description
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Syntax

Click **Show Syntax** to view the EPL syntax of conditions, statements, and debugging parameters. It also provides a warning when the syntax is invalid. For more information, see [Rule Syntax Dialog](#).

Build a Statement Dialog

The Build a Statement dialog allows you to construct a condition statement when creating a new Rule Builder rule.

What do you want to do?




Role	I want to ...	Show me how
Content Expert	Configure a rule statement.	Step 2. Build a Rule Statement
Content Expert	Add conditions to the rule.	Step 3. Add Conditions to a Rule Statement

Related Topics

- [Add a Rule Builder Rule](#)

Quick Look

To access the Build a Statement dialog:

1. Go to  **(Configure)** > **ESA Rules**.
The Configure ESA Rules view is displayed with the Rules tab open.
2. In the **Rule Library** toolbar, select  > **Rule Builder**.
A New Rule tab is displayed..
3. In the **Conditions** section, click  .
The Build a Statement dialog is displayed.

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name * 5 Failed Logins



if all conditions are met + -

Key	Operator	Value	Ignore Case?	Array?
<input checked="" type="checkbox"/> event.ec_outcome	is	Failure	<input type="checkbox"/>	<input type="checkbox"/>

Cancel Save

The following table describes the parameters in the Build a Statement dialog.

Parameter	Description
Name	Purpose of the statement.
Select	Conditions the rule requires. There are two options: <ul style="list-style-type: none"> • If all conditions are met • If any of these conditions are met
Key	Key for ESA to check in the rule statement.

Parameter	Description
Operator	Relationship between the meta key and value for the key: <ul style="list-style-type: none"> • is • is not • is not null • is greater than (>) • is greater than or equal to (>=) • is less than (<) • is less than or equal to (<=) • is one of (For array type meta) • is not one of (For array type meta) • contains • not contains • begins with • ends with
Value	Value for ESA to look for in the key.
Ignore Case?	This field is designed for use with string and array of string values. By choosing the Ignore Case field, the query will treat all string text as a lowercase value. This ensures that a rule that searches for the user named Johnson would trigger if the event contains "johnson," "JOHNSON," or "JoHnSoN."
Array?	Choice to indicate if contents of Value field represent one value or multiple values: <ul style="list-style-type: none"> • Select the box to indicate multiple values. • Clear the box to indicate one value.
	Add a statement. You can add a meta condition, whitelist condition, or blacklist condition.
	Delete selected statement.
Save	Add statement to the Conditions section of the Rule Builder tab.

The following table shows the operators you can use in the Rule Builder:

Operator	Required Value	Usage	Example	Meaning
is	Singular string value	The meta key is equal to the <i>value</i> field.	<i>user_dst</i> is John Doe.	<i>user_dst</i> is equal to the string " <i>John Doe</i> ".

Operator	Required Value	Usage	Example	Meaning
is	Array string value	The meta key is equal to one of the elements of the <i>value</i> field.	<i>user_dst</i> is John, Doe, Smith.	<i>user_dst</i> is equal either to the string "John" or to the string "Doe" or to the string "Smith" (Note, the spaces are stripped.).
is not	Singular string value	The meta key is not equal to the <i>value</i> field.	<i>size</i> is not 200.	<i>size</i> is not equal to the number 200 (<i>size</i> is a numeric value).
is not	Array string value	The meta key is not equal to any of the elements of the <i>value</i> field.	<i>size</i> is not 200, 300, 400.	<i>size</i> is equal neither to 200 nor to 300 nor to 400.
is not null	N/A (looks for any value)	The meta key value is not null.	<i>user_dst</i> is not null.	<i>user_dst</i> is a meta that contains a value.
is greater than (>)	Number	The numeric value of the meta key is greater than the number in the <i>value</i> field.	<i>payload</i> is greater than 7000.	<i>payload</i> is a numeric value that is greater than 7000.
is greater than or equal to (>=)	Number	The numeric value of the meta key is greater than or equal to the number in the <i>value</i> field.	<i>payload</i> is greater than or equal to 7000.	<i>payload</i> is a numeric value that is greater than or equal to 7000.
is less than (<)	Number	The numeric value of the meta key is less than the number in the <i>value</i> field.	<i>ip_dstport</i> is less than 1024.	<i>ip_dstport</i> is a numeric value that is less than the numeric value 1024.
is less than or equal to (<=)	Number	The numeric value of the meta key is less than or equal to the number in the <i>value</i> field.	<i>ip_dstport</i> is less than or equal to 1024.	<i>ip_dstport</i> is a numeric value that is less than or equal to numeric value 1024.
is one of	Array string value	The meta key is one of the array string values in the <i>value</i> field.	<i>alias_host</i> is one of Facebook, UTube, Instagram.	<i>alias_host</i> is one of the array string values <i>Facebook</i> , <i>UTube</i> , <i>Instagram</i> .
is not one of	Array string value	The meta key is not one of the array string values in the <i>value</i> field.	<i>alias_host</i> is not one of Facebook, UTube, Instagram.	<i>alias_host</i> is not one of the array string values <i>Facebook</i> , <i>UTube</i> , <i>Instagram</i> .

Operator	Required Value	Usage	Example	Meaning
contains	String	The <i>value</i> field is a substring of the meta key. (This operator is only available for a string-valued meta key).	<i>ec_outcome</i> contains failure.	<i>ec_outcome</i> is a string that contains the substring "failure".
not contains	String	The <i>value</i> field is not a substring of the meta key (This operator is only available for a string-valued meta key).	<i>ec_outcome</i> not contains failure.	<i>ec_outcome</i> is a string that does not contain the substring "failure".
begins with	String	The <i>value</i> field is the beginning of the meta key (This operator is only available for a string-valued meta key).	<i>ip_dst</i> begins with 127.0.	<i>ip_dst</i> is a string that starts with "127.0".
ends with	String	The <i>value</i> field is the end of the meta key (This operator is only available for a string-valued meta key).	<i>user_dst</i> ends with son.	<i>user_dst</i> is a string that ends in "son".

Note: Terms in *bold italics* are Meta that may not exist in all customer environments.

Advanced EPL Rule Tab

The Advanced EPL Rule tab enables you to define rule criteria with an Event Processing Language (EPL) query.

What do you want to do?



Role	I want to ...	Show me how
Content Expert	Define an Advanced EPL rule.	Add an Advanced EPL Rule
Content Expert	Test the Advanced EPL rule logic.	Validate an Advanced EPL Rule
Content Expert	See examples of an Advanced EPL Rule.	Example Advanced EPL Rules
Content Expert	See best practices for writing Advanced EPL Rules.	ESA Rule Writing Best Practices

Related Topics

- [Add a Rule Builder Rule](#)
- [Enrichment Sources](#)

Quick Look

To access the Advanced EPL Rule tab:

1. Go to  (Configure) > **ESA Rules**.
The Configure view is displayed with the Rules tab open by default.
2. In the **Rule Library** toolbar, select  > **Advanced EPL**.
The Advanced EPL Rule tab is displayed.

The following figure shows the Advanced EPL Rule tab.

Advanced EPL
Write a rule in Event Processing Language.

Rule Name *

Description

Trial Rule

Memory Threshold MB

Alert

Severity *

Query *

```
@audit('stream')
@RSAAlert
SELECT * FROM Event(((ip_src NOT IN ('1.1.1.1', '2.2.2.2'))));
```

Notifications [Global Notifications](#)

Output	Notification	Notification Server	Template
No parameters to edit.			

Output Suppression of every minutes

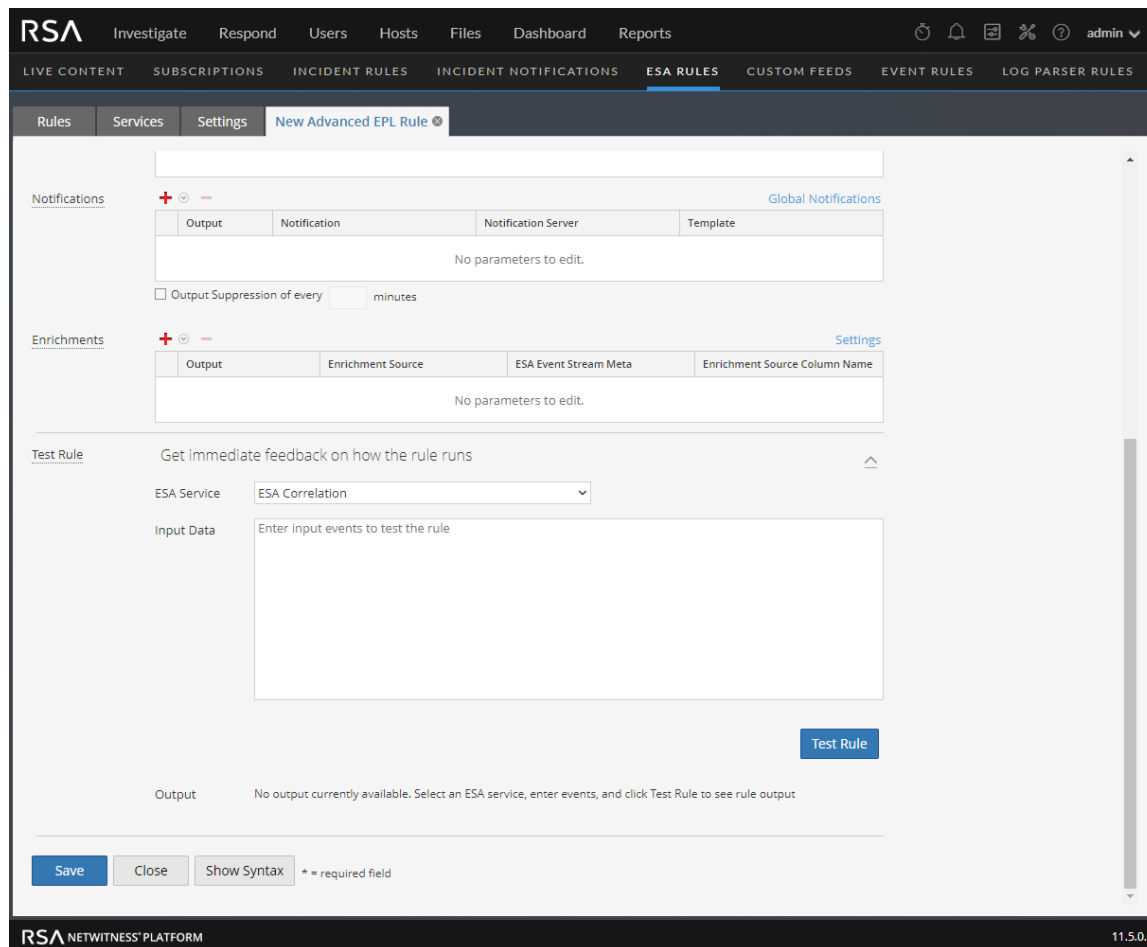
Enrichments [Settings](#)

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
No parameters to edit.			


Test Rule [Get immediate feedback on how the rule runs](#)

RSA NETWITNESS PLATFORM 11.5.0.0

The following figure shows the Advanced EPL Rule tab scrolled down with the Test Rule section in view.



The following table lists the parameters in the Advanced EPL Rule tab.

Parameters	Description
Rule Name	Purpose of the ESA rule.
Description	Summary of what the ESA rule detects.
Trial Rule	Deployment mode to see if the rule runs efficiently.
Memory Threshold	(This option applies to version 11.5 and later.) The maximum memory usage allowed for this rule in MB. Add Memory Thresholds to ESA rules that use memory. For example, if a rule contains windows or pattern matching, configure a memory threshold for that rule. If the configured memory threshold is exceeded, it gets disabled individually and an error is displayed for that rule on the  (Configure) > ESA Rules > Services tab. New rules default to a 100 MB memory threshold. Rules that existed before version 11.5 do not have a default value and a memory threshold is not set.

Parameters	Description
Alert	(This option applies to version 11.3 and Later.) When selected, the alert is sent to Respond. If the checkbox is cleared, an alert will not be sent to Respond. To turn alerts on or off for ALL rules, see the <i>ESA Configuration Guide</i> .
Severity	Threat level of alert triggered by the rule.
Query	EPL query that defines rule criteria.



Notifications Section

In the Notifications section, you can choose how to be notified when ESA generates an alert for the rule. For more information on the alert notifications, see [Add Notification Method to a Rule](#).

The following figure shows the Notifications section.

Output	Notification	Notification Server	Template
<input checked="" type="checkbox"/> SYSLOG	Local_SysLog	localhost-514	Default Syslog Template

Output Suppression of every minutes

Parameter	Description
	To add an alert notification type.
	To delete the selected alert notification type.
Output	Alert notification type. Options are: <ul style="list-style-type: none"> Email SNMP (This option is not supported in NetWitness version 11.3 and later.) Syslog Script
Notification	Name of previously configured output, such as an email distribution list.
Notification Server	Name of server that sends the output.
Template	Name of template for the alert notification.
Output Suppression of every	Option to specify alert frequency.
Minutes	Alert frequency in minutes.

Enrichments Section

In the Enrichments section, you can add a data enrichment source to a rule.

For more information on the enrichments, see [Add an Enrichment to a Rule](#).

The following figure shows the Enrichments section.

Output	Enrichment Source	ESA Event Stream Meta	Enrichment Source Column Name
<input type="checkbox"/> In-Memory Table	Select Enrichment Source	Enter Meta	Enter Column Name
<input checked="" type="checkbox"/> GeoIP	Select Enrichment Source	Enter Meta	ipv4

Parameter	Description
	To add an enrichment.
	To delete the selected enrichment.
Output	Enrichment source type. Options are: <ul style="list-style-type: none"> In-Memory Table (Ad hoc only - Recurring In-Memory Tables are no longer supported in version 11.3 and later.) GeoIP
Enrichment Source	Name of previously configured enrichment source, such as a .CSV filename for an In-Memory Table.
ESA Event Stream Meta	ESA meta key whose value will be used as one operand of join condition.
Enrichment Source Column Name	Enrichment source column name whose value will be used as the other operand of the join condition.

Test Rule Section

Note: The Test Rule section is available in NetWitness Platform 11.5 and later.

In the Test Rule section, you can validate your ESA rule to determine if the rule logic is working as expected before deploying the rule.

Test Rule Get immediate feedback on how the rule runs ^

ESA Service:

Input Data:

```
"Authentication"
],
"ec.outcome": [
  "Failure"
],
"reference.id": [
  "605004"
],
"event.desc": [
  "Login denied"
]
```

Test Rule

Output: Test complete

- ✓ Rule successfully validated
- ✓ Provided input is valid
- ✓ Test ran successfully

Engine Stats

Engine Version	Events Offered	Offered Rate	Runtime Errors
8.4.0	9	0	-

Rule Stats

Deployed	Statements Fired	Alerts Fired	Events in Memory	Memory Usage	CPU %	Events Matched	Alerted Events	Runtime Errors	Debug Logs
✓	0	9	0	0	100	9	Details...	-	Details...

Field	Description
ESA Service	Select the ESA Correlation service to process the rule.
Input Data	Enter the input events to test the rule. You can download the events from the Investigate view in JSON format, copy the events, and paste them in this field.
Output Data	After you select an ESA Correlation service, input data, and click the Test Rule button, you can view the output of the rule here and verify that the rule is working according to your requirements. You can view the alerts in the output, but this test does not send any alert notifications. If you want to view all of the debug information for the test, include an <code>@Audit('stream')</code> annotation to your rule query.

The following table describes the test rule output **Engine Stats**.

Field	Description
Engine Version	Esper version running on the ESA service

Field	Description
Events Offered	Number of events processed by the ESA service since the last service start
Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the ESA rule deployment.

The following table describes the test rule output **Rule Stats**.

Field	Description
Deployed	A green checkmark indicates that the rule is deployed on the selected ESA service.
Statements Fired	The number of statements that fired the alerts
Alerts Fired	The number of alerts generated from the test data
Events in Memory	The number of events placed in memory by the rule
Memory Usage	The total amount of memory used by the rule
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each.
Events Matched	The number of events that matched the rule
Alerted Events	If applicable, this field can contain a link to events that caused an alert.
Runtime Errors	If applicable, this field can contain a link to runtime error messages related to the rule.
Debug Logs	This field contains a link to Esper debug (audit) logs.

Syntax







Click **Show Syntax** to view the EPL syntax of conditions, statements, and debugging parameters. It also provides a warning when the syntax is invalid. For more information, see [Rule Syntax Dialog](#).

Rule Syntax Dialog


This topic describes the features of the Rule Syntax dialog. The Rule Syntax dialog displays the EPL syntax of conditions, statements, and debugging parameters, and provides a warning when the syntax is invalid.

Quick Look

To access this dialog:

1. Go to  **(Configure) > ESA Rules**.
2. In the **Rule Library** view, do one of the following:
 - a. Click   and select **Advanced EPL** or **Rule Builder**.
 - b. Double-click an existing rule.
 - c. Select an existing rule and click  in the **Rule Library** toolbar.
 - d. In the row of an existing rule, select   > **Edit**.
The new or existing rule is displayed in a new tab, available to edit.
3. Click **Show Syntax** at the bottom of the tab.


The following figure shows an example of the Rule Syntax dialog showing a valid rule.



The following table describes the Rule Syntax dialog parameters.

Parameters	Description
Rule is valid or Validation error in rule	Indicates whether the rule syntax is valid or needs to be changed.
Rule Name	Displays the name of the rule.
Text	Displays the EPL syntax of conditions, statements, and debugging parameters if the rule is valid.

Deployment Panel

ESA rule deployments map rules from your rule library to the appropriate ESA Services and data sources. The Deployment panel ( (Configure) > ESA Rules > Rules tab) enables you to create and configure ESA rule deployments that specify:

- ESA Services
- Data Sources (This is available in NetWitness version 11.3 and later.)
- ESA Rules

When you are ready to start aggregating data and generating alerts from an ESA rule deployment, you deploy the ESA rule deployment to activate it.

Note: An ESA rule deployment can have only one ESA service. You can, however, use the same ESA service in multiple deployments. In NetWitness Platform version 11.2 and earlier, the ESA service is the Event Stream Analysis service. In version 11.3 and later, it is the ESA Correlation service.

What do you want to do?


Role	I want to ...	Show me how
Content Expert	Add an ESA rule deployment.	ESA Rule Deployment Steps
Content Expert	Manage deployments.	Additional ESA Rule Deployment Procedures

Related Topics

- [View Stats for an ESA Service](#)

Quick Look



The following figure shows the Deployment panel.

- 1 Shows the options panel. Click the deployment in the options panel to view the deployment in the Deployment panel on the right.
- 2 Shows the Deployment panel.
- 3 Shows the ESA rule deployment name.
- 4 The ESA Services section shows the ESA Correlation service that processes the ESA rules and creates alerts.
- 5 The Data Sources section shows the data sources, such as Concentrators, that provide the data for the deployment.
- 6 The deployment options enable you to view the changes to the deployment (Show Updates) and deploy the ESA rule deployment on ESA (Deploy Now). The  (Configure) > ESA Rules > Services tab shows the status of the deployment after it is deployed.
- 7 The ESA Rules section shows the rules that are used to trigger alerts in the deployment.
- 8 The Data Source Filter (Optional) section enables you to add a filter query to forward only the data relevant to this deployment to ESA. The data source filter is for advanced users familiar with Decoder application rules.

ESA Services Section

In the ESA Services section, you can manage each ESA service in the deployment.

The following table describes the actions you can perform in the ESA Services section.

Task	Description
	Adds an ESA service to the deployment.
	Removes the selected ESA service from the deployment.

The following table describes the columns in the ESA Services section.




Title	Description
Status	Indicates if the deployment status is Added , Deployed , Updated , or Failed .
Name	Name of the ESA service.
Address	IP address of the host where the ESA service is installed.
Version	Version of the ESA service.
Last Deployment Date	The date and time when the ESA service was last deployed.

Data Sources Section

Note: This option is available in NetWitness Platform version 11.3 and later.

In the Data Sources section, you can select one or more data sources, such as Concentrators, to use for your selected ESA Service.

The following table describes the actions you can perform in the Data Sources section.

Task	Description
	Adds a data source for the selected ESA service to the deployment.
	Removes a data source for the selected ESA service from the deployment.
	<p>(This option is available in NetWitness Platform version 11.3.0.2 and later.) Enables you to change the configuration of a data source in an ESA rule deployment. You can change the data source password, SSL, port, and compression settings. When a data source password changes, it is important to change the password on the data source so that ESA can continue to communicate with the data source.</p> <p>Note: If you make any ESA service, data source, or ESA rule changes to an ESA rule deployment, you need to redeploy the deployment. For example, if you change the configuration of a data source in an ESA rule deployment, you must redeploy all the ESA rule deployments that contain that data source. When you set the compression level for a Concentrator on ESA, it sets the same compression level for that Concentrator for ESA Correlation Rules.</p>

The following table describes the columns in the Data Sources section.

Title	Description
(Status)	Shows the status of the data source. A solid colored green circle indicates a running service and a white circle indicates a stopped service.
Name	Shows the name of the data sources used by the selected ESA service. You can specify the data sources separately for each ESA rule deployment.
Type	Shows the type of the data sources. Data sources can be Concentrators or Decoders. It is important that you choose data sources that have the appropriate data for the rules in the deployment. For example, if you have NetWitness Endpoint and you want to deploy the Endpoint Risk Scoring Rules Bundle, you must choose endpoint data sources.

Note: You can add a Log Decoder as a data source for ESA, but it is better to add a Concentrator to take advantage of undivided aggregation as the Decoder may have other processes aggregating from it.

Position Tracking Information

The ESA Correlation service continuously streams data from the data sources like decoders (log and network), and concentrators. ESA retrieves events from the data sources, and applies rules to generate alerts to detect malicious activities. When you deploy a data source, ESA starts processing information from the latest available session, by default. Position Tracking Information enables you to visualize the progress of the sessions that ESA has processed, and provides information on the session IDs and the time/date when the events were processed.

Edit Position Tracking Information enables you to:


- Visualize the number of sessions that a particular ESA data source has already analyzed, review the number of sessions ESA would process after you edit the position tracking, and plan your work.
- Edit the tracking position information based on:
 - Date and Time (Collection Time)
 - Session ID
- Edit position tracking for multiple data sources before you deploy them.
- Calculate the number of sessions that the ESA Correlation Service is scheduled to process for a particular data source to either process, reprocess, or skip sessions with respect to the current position of the data source.

Note: The Edit Position tracking feature with the `Date and Time` option works based on the profile time settings in the NetWitness UI. This time-zone based time from the UI is converted to UTC, and is sent to the core, to retrieve the corresponding session ID for that time stamp. Example: If the UI follows IST, the UI converts it to UTC and sends it to the core. The session ID is fetched for the specific UTC time stamp, and set to position tracking at deployment.

Editing Position Tracking Information

The following figure shows the Edit Position Tracking window.

To edit position tracking information:

1. Select the specific data source from the **Data Source** menu.
2. Click  .

The **Edit Service** window is displayed.

3. Select the **Trusted Authentication** check-box, or enter your administrative credentials (username and password).

4. In the **Position Tracking Information** menu, click the **Edit Tracking** check-box to select it.

a. If you want to edit the position tracking information based on date and time stamp:

In the **Go To** text field, select **Date and Time** and enter the date and time.

The ESA Correlation service starts processing the events from the date and time that you entered.

b. If you want to edit the position tracking information, based on the session ID:

In the **Go To** text field, select **Session ID** and enter the session ID.


The ESA Correlation service starts processing the events from the session ID that you entered.

5. Click **Calculate Sessions** to calculate the number of sessions that will be processed with respect to the existing position of the data source, if any.

6. Click **OK**.

7. Click **Deploy Now**.

The tracking position information will be deployed to the ESA Correlation service, only when the deployment is successfully completed.

Note: After you deploy, in the **Data Sources** menu, click on , to view the edit tracking information. The default time-out associated with this information is 1 minute.

Use Case Scenario

This section provides information about how you can use position tracking information in a real world scenario.

If you have deployed a data source with a total of 72 sessions that ESA has already processed, and if you want to start processing the events from the beginning, or go back with respect to the time or sessions ID.


1. Click **Edit Tracking**.

2. Enter 1 in the **Session ID** text box.

3. Click **Calculate Sessions**.

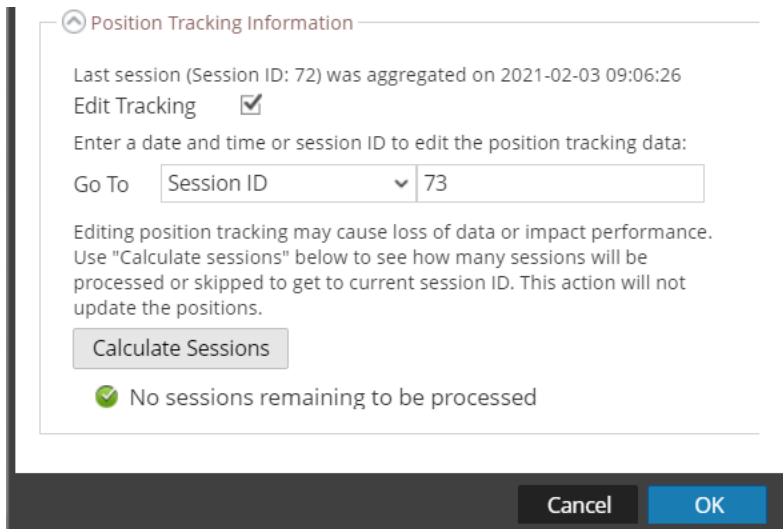
All the 72 sessions will be reprocessed.

The following image shows the use case scenario.



4. Enter a future Session ID.

The **No sessions remaining to be processed** message is displayed.



Note: Editing the tracking information is optional. If you add a new data source to an existing ESA deployment, and you do not edit the tracking information, ESA follows the default behavior to process events.

Importing Position Tracking Information

You can migrate the settings of position tracking for one or more data sources at the same time from an existing deployment, using the import function.

To import position tracking information from an existing deployment:

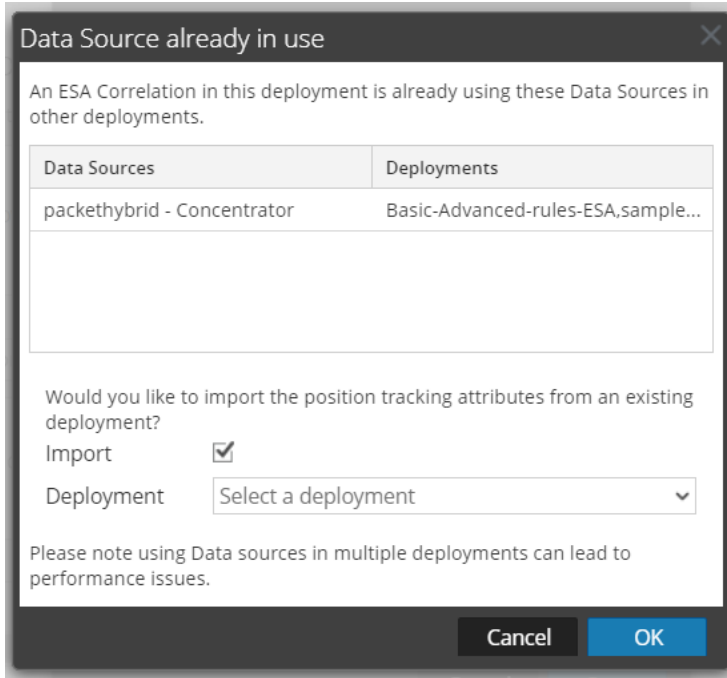
1. Go to **Configure > ESA Rules > Deployment** tab.
2. Create a new deployment. See [ESA Rule Deployment Steps](#)
3. Add ESA Services. See [Step 2. Add an ESA Service](#)

The Data Sources panel will be enabled.

4. Add Data Sources. See [Step 3. Add Data Sources](#)
5. In the Data Sources tab, click **+** to add a configured data source. The Available Configured Data Sources window appears.

6. Select the Data Sources and click **Save**. The Data Sources window displays the data sources that are already in use by the other deployments.

The following figure shows the Import functionality.



7. Select the **Import** check-box.

8. In the **Deployment** dropdown box, select the deployment from which you want to import the position tracking attributes from an existing deployment. The dropdown menu lists all of the deployments with which data sources are shared. Review the list of data sources in the **Data Sources** menu, and click **OK**.

9. Add an ESA Rule. See [Step 4. Add and Deploy Rules](#)

10. In the Data Sources tab, click **Deploy Now**.

The deployment retrieves the position tracking information from the selected pre-existing deployment, and applies it to the new deployment.

Note: Ensure that you do not delete the deployment from where you imported the position tracking information, before you perform **Deploy Now** on the newly created or edited deployment. The import position tracking information function is disabled by default. Sharing data sources between multiple deployments can lead to performance issues.

Deployment Options

There are two deployment options below the Data Sources section. These options apply to the entire ESA rule deployment.




The following table describes these deployment options.

Task	Description
Show Updates	Enables you to view a history of updates to the deployment.
Deploy Now	Activates the ESA rule deployment. The selected ESA service starts aggregating data and generating alerts using the specified ESA rules in the deployment. You need to add ESA Rules to the deployment before deploying the ESA rule deployment.



ESA Rules Section

In the ESA Rules section, you manage rules in the deployment. This section lists all rules that are currently in the deployment.

The following table describes the actions you can perform in the ESA Rules section.

Task	Description
	Opens the Deploy ESA Rules dialog, where you can select a rule.
	Removes the selected ESA rules from the deployment.
	Filters the list of rules.
<input type="text" value="Filter"/>	Enables you to search for a rule.

The following table describes the columns in the ESA Rules section.

Title	Description
Status	<p>Indicates the rule status:</p> <ul style="list-style-type: none"> • Deployed - the rule is deployed. • Updated - the rule has been updated since the last deployment. • Added - the rule has been added since the last deployment. • Disabled - the rule is disabled due to an error in the rule or an error during the deployment of the rule. <p>In NetWitness Platform version 11.3.0.2 and later, if a disabled rule has an error message, it shows  in the Status field. Hover over the rule to view the error message tooltip.</p> 
Rule Name	Describes the purpose of the ESA rule.
Trial Rule	Indicates whether the rule is Deployment mode to see if the rule runs efficiently.
Severity	Shows the threat level of alert triggered by the rule.
Type	Shows the type of the ESA rule. For more information, see ESA Rule Types .

Title	Description
Email, SNMP, Syslog, Script	Indicates which notification types are used for alerts generated by the rules. (ESA SNMP notifications are not supported in NetWitness version 11.3 and later.)
Last Modified	Shows the date and time when the ESA rule was last modified.

Data Source Filter (Optional) Section


Note: This option is available in NetWitness version 11.5 and later.

The data source filter is optional. If you have a medium to large NetWitness Platform deployment and you have high throughputs, you can add a filter query to forward only the data relevant to this deployment to ESA. This helps to filter out certain types of traffic that does not add value to the analysis of the data in the ESA rule deployment.

Caution: The data source filter is intended for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Using a data source filter can be performance intensive for data aggregation. A filter slows the event aggregation rate, but when you are filtering a large amount of traffic, it can have performance benefits on ESA Correlation server. However, if you use a complex filter and do not filter a large amount of traffic, the event aggregation rate may be lower than expected.

IMPORTANT: If an application rule linked to a data source filter is modified on a Decoder, the filter must be removed, added again, and redeployed. The changes take effect on ESA after the deployment is redeployed.



The following table describes the columns in the Data Source Filter (Optional) section.

Title	Description
<input type="checkbox"/>	Enables you to select and remove a data source filter.
Status	Indicates the data source filter status: <ul style="list-style-type: none"> Added - the filter is added to the deployment, but it is not yet deployed. Deployed - the filter is actively streaming only the relevant data as defined in the filter query.

Title	Description
Filter Query	Shows the query that contains the application rules. Each application rule query is separated by an "or" condition. You can use either the simple or advanced option to create the data source filter. You can only have one data source filter per deployment for your selected data sources.
Last Modified	Shows the date that the filter was added to the ESA rule deployment.

Create Data Source Filter Dialog - Simple

When you create the data source filter, you select application rules to be included in the filter query. The application rules that you select must be enabled on the Decoders that feed the data sources in this deployment. ESA Correlation uses the filtered event data to process the ESA rules.

Caution: The data source filter is for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Simple Advanced

Select the application rules that are relevant to the data required for this deployment. Any application rule that is not included in this filter does not get aggregated by ESA.
 The application rules that you select must be enabled on the Decoders that feed the data sources in this deployment. ESA Correlation uses the filtered event data to process the ESA rules.

Application Rules for Selected Deployment

Filter

<input type="checkbox"/>	Application Rule Name	Alert Field	Present On	Absent On
<input type="checkbox"/>	account:deleted	alert.id	2	0
<input type="checkbox"/>	account:modified	alert.id	2	0
<input type="checkbox"/>	config:config-changes	alert.id	2	0
<input type="checkbox"/>	config:fw-config-changes	alert.id	2	0
<input type="checkbox"/>	config:router-change	alert.id	2	0
<input type="checkbox"/>	encryption:failures	alert.id	2	0
<input type="checkbox"/>	encryption:key-gen-and-changes	alert.id	2	0
<input type="checkbox"/>	fw:categories	alert.id	2	0
<input type="checkbox"/>	fw:inbound-network-traffic	action	2	0
<input type="checkbox"/>	fw:outbound-network-traffic	action	2	0
<input type="checkbox"/>	host:windows:account-disabled	alert.id	2	0
<input type="checkbox"/>	intrusion:all-activity	alert.id	2	0
<input type="checkbox"/>	ssl-lh2	alert.id	1	1
<input type="checkbox"/>	ssl-lh1	alert.id	1	1

rsa-nw-11.4-lh1 - Log Decoder

Cancel Save

The following table describes the columns in the Create Data Source Filter dialog.

Title	Description
Filter	Enables you to search for a rule. For example, you can type "account" to search for application rules that contain that word.
Application Rule Name	Shows the names of the application rules on the Decoders mapped to the data sources in this deployment.
Alert Field	Shows the meta key used in the alert.
Present On	Shows the number of Decoders mapped to the data sources in this deployment that <i>have</i> the listed application rule. If the rule is not present on all of the Decoders, you can hover over the ⓘ icon to view the names of the Decoders that contain the rule.
Absent On	Shows the number of Decoders mapped to the data sources in this deployment that <i>do not have</i> the listed application rule. If a rule is not present on all of the Decoders, you can hover over the ⓘ icon to view the names of the Decoders that do not contain the rule.

For more information, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.

Create Data Source Filter - Advanced

If necessary, you can use the advanced filter instead of the simple filter to add your data source filter query directly. The individual application rule queries must be separated by an "or" condition. For more information on creating and writing Decoder rules, see "Configure Application Rules" in the *Decoder and Log Decoder Configuration Guide*.

Caution: The data source filter is for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Create Data Source Filter

Simple Advanced

A data source filter can improve performance by only aggregating the data that is relevant to an ESA rule deployment. For the filter to work, application rules must be enabled on the mapped Decoders to apply the meta keys required for this filter. After saving the filter, it must be deployed for it to take effect.

Caution: The data source filter is for advanced users familiar with Decoder application rules. Improper filtering can cause the required data to not be forwarded to and analyzed by ESA.

Data Source Filter

```
select * where ((alert.id = 'account:logon-success') or (alert.id = 'account:logon-failure'))
```

Example :
*select * where service = 443*

Cancel Save

Deploy ESA Services Dialog

The Deploy ESA Services dialog displays all ESA services available to be added to an ESA rule deployment.

What do you want to do?



Role	I want to ...	Show me how
Content Expert	Configure an ESA rule deployment.	ESA Rule Deployment Steps
Content Expert	Add a service.	ESA Rule Deployment Steps
Content Expert	Add data sources. (This is available in 11.3 and later.)	ESA Rule Deployment Steps
Content Expert	Add and deploy rules.	ESA Rule Deployment Steps

Related Topics

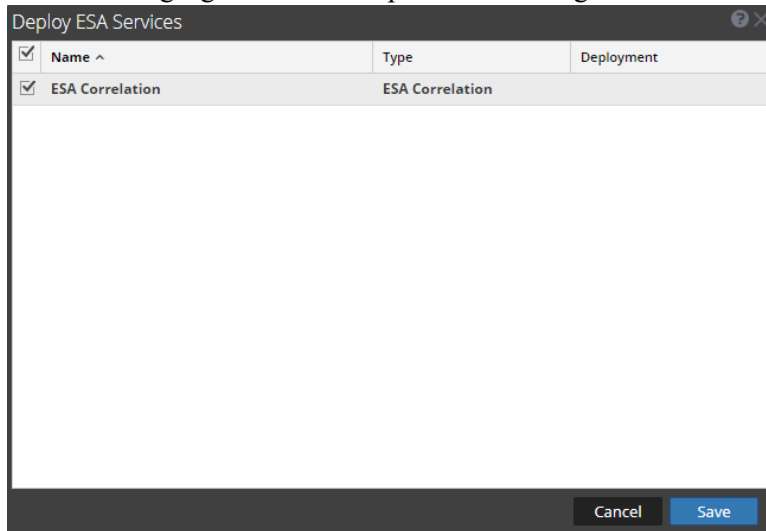
- [Additional ESA Rule Deployment Procedures](#)
- [View Stats for an ESA Service](#)

Quick Look

To access this dialog:

1. Go to  (**Configure**) > **ESA Rules**.
The Rules tab opens by default.
2. In the options panel, under the **Deployment** section, select or add a deployment.
3. In the **ESA Services** panel, click  .
The Deploy ESA Services dialog is displayed.

The following figure is an example of this dialog.



The following table describes the parameters of the Deploy ESA Services dialog.

Parameters	Description
Name	Displays the name of configured ESA services.
Deployment	Displays the ESA rule deployments to which the service has already been added.

Deploy ESA Rules Dialog

The Deploy ESA Rules dialog enables you to filter and select rules to deploy to an ESA service.

What do you want to do?





Role	I want to ...	Show me how
Content Expert	Configure an ESA rule deployment.	Step 1. Add an ESA Rule Deployment
Content Expert	Deploy a rule	Step 4. Add and Deploy Rules

Related Topics

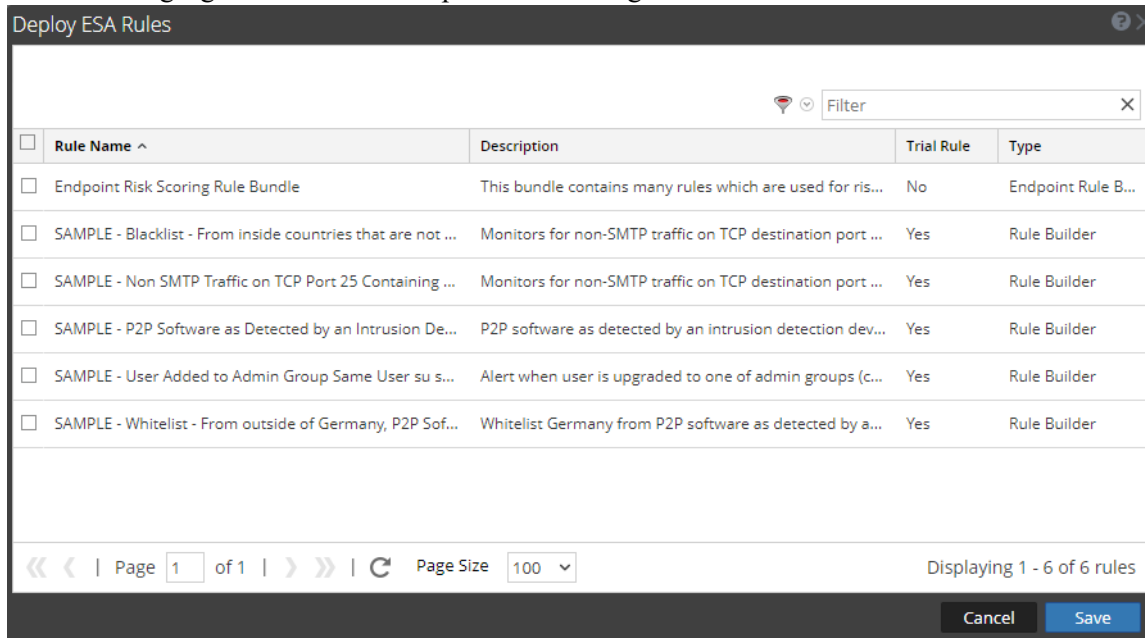
- [Additional ESA Rule Deployment Procedures](#)

Quick Look


To access this dialog:

1. Go to  **(Configure) > ESA Rules**.
The Rules tab opens by default.
2. In the options panel, under the **Deployment** section, select or add a new deployment by clicking   > **Add**.
3. If you add a new deployment, type the name of the deployment in the box in the options panel.
4. In the **ESA Rules** panel, click  .
The Deploy ESA Rules dialog is displayed.


The following figure shows an example of this dialog.



The following table describes the parameters of the Deploy ESA Rules dialog.

Parameters	Description
	Filters the list of rules based on severity and type. The text box beside this icon filters based on rule name.
Rule Name	Displays the name of the rule.
Description	Describes the rule.
Trial Rule	Indicates whether or not the rule is a trial rule.
Type	Indicates the type of rule: RSA Live ESA, Advanced EPL, or Rule Builder.

Updates to the Deployment Dialog

The Updates to the Deployment dialog displays changes made to the deployment, such as adding a rule or service. Deployment updates are indicated by the update icon () next to the name of the deployment in the Rules tab options panel.

What do you want to do?


Role	I want to ...	Show me how
Content Expert	Deploy rules to run on ESA.	ESA Rule Deployment Steps
Content Expert	Edit or delete an ESA rule deployment.	Edit the ESA Rule Deployment Name or Delete a Deployment
Content Expert	View deployment updates.	Show Updates to an ESA Rule Deployment

Related Topics

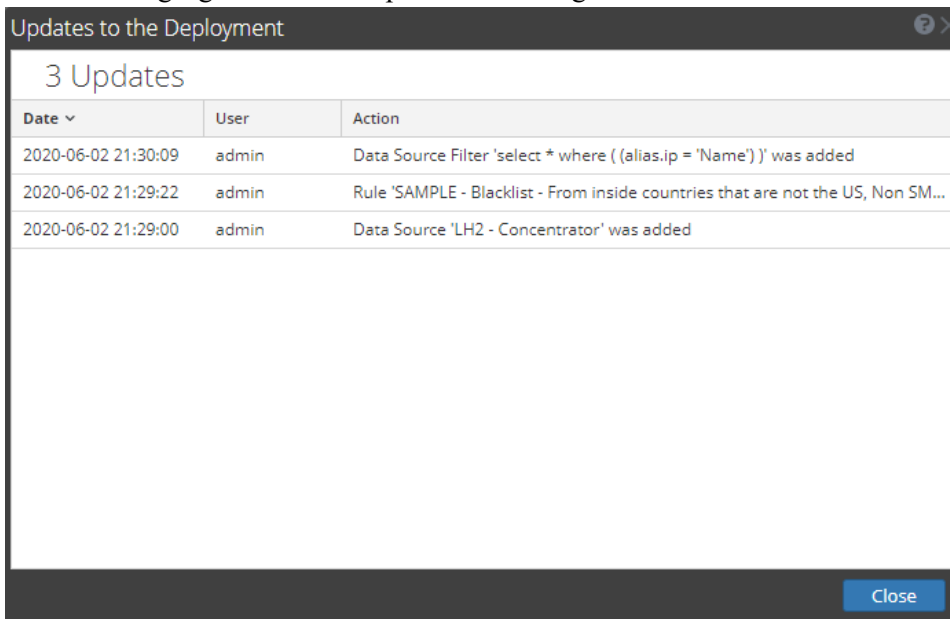
- [Replace an ESA Service in an ESA Rule Deployment](#)
- [Add or Remove a Data Source](#)
- [Edit or Delete a Rule in a Deployment](#)

Quick Look

To access this dialog:

1. Go to  (**Configure**) > **ESA Rules**.
The Rules tab opens by default.
2. In the options panel, under the **Deployments** section, select or add a deployment.
3. In the **Deployment** panel, click **Show Updates**.
The Updates to the Deployment dialog is displayed.


The following figure is an example of this dialog.



The Updates to the Deployment dialog displays the number of updates at the top of the dialog. The following table describes the parameters of this dialog.

Parameters	Description
Date	Displays the day and time of the update.
User	Displays the user who made the update.
Action	Describes the update.

Services Tab

This topic provides an overview of the  (Configure) > **ESA Rules** > **Services** tab. The Services tab shows the status of the deployments on each ESA service.

What do you want to do?


Role	I want to ...	Show me how
Content Expert	Troubleshoot Services Tab.	Troubleshoot ESA
Content Expert	View deployment Stats for an ESA Service.	View Stats for an ESA Service

Related Topics

- [View a Summary of Alerts](#)

Quick Look

The following figure shows the Services tab:



The screenshot shows the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. The main navigation bar is currently on 'ESA RULES'. The sub-navigation bar has 'Rules', 'Services', and 'Settings'. The 'Services' tab is active, showing 'esaSecondary - ESA Correlation'. There are two deployment tabs: 'Deployment A' and 'Deployment B'. The 'Deployment A' tab is selected and displays the following statistics:

Engine Stats		Rule Stats		Alert Stats	
Esper Version	8.4.0	Rules Enabled	75	Notifications	0
Time	2020-05-05T19:44:43	Rules Disabled	0	Message Bus	1
Events Offered	74750	Events Matched	29		
Offered Rate	0 per second / 8,557 max				
Status	Active				

Below the statistics is a section for 'Deployed Rule States' with a table listing various rules. The table has columns for 'Enabled', 'Name', 'Rule Type', 'Trial Rule', 'Last Detected', 'Events Matched', 'Memory Usage', and 'CPU %'. The rules listed include 'AWS Critical VM Modified', 'AWS Permissions Modified Followed By Instan...', 'Account Added to Administrators Group and ...', 'Account Removals From Protected Groups on ...', 'Aggressive Internal Database Scan', 'Aggressive Internal NetBIOS Scan', 'Aggressive Internal Web Portal Scan', 'BYOD Mobile Web Agent Detected', 'Backdoor Activity Detected', and 'Cerber Ransomware'. At the bottom, there is a pagination control showing 'Page 1 of 1' and 'Page Size 100', along with the text 'Displaying 1 - 75 of 75'.


(This option is available in NetWitness version 11.3 and later.) If an ESA Correlation service has multiple deployments, under the service name, you will see a tab for each deployment. In the above example, there are two deployment tabs, Deployment A and Deployment B. Each tab displays information specific to that deployment.

The Services tab has the following sections:

- ESA Services panel (on the left)
- General Stats panel (top right)
- Deployed Rule Stats panel (bottom right)

ESA Services Panel

The ESA Services panel lists the name of each ESA service added to NetWitness.



General Stats Panel

The General Stats panel provides information on the Esper engine, rules, and alerts.

The General Stats panel contains the following sections:

- Engine Stats
- Rule Stats
- Alert Stats

The following figure shows the General Stats panel.

The screenshot shows a window titled 'esaSecondary - ESA Correlation'. There are two tabs: 'Deployment A' (selected) and 'Deployment B'. The content is organized into three columns:

Engine Stats		Rule Stats		Alert Stats	
Esper Version	8.4.0	Rules Enabled	75	Notifications	0
Time	2020-05-05T19:44:43	Rules Disabled	0	Message Bus	1
Events Offered	74750	Events Matched	29		
Offered Rate	0 per second / 8,557 max				
Status	Active				

The following table lists and describes the parameters in each section.

Sections	Parameter	Description
Engine Stats	Esper Version	Esper version running on the ESA service
	Time	Time when the last event was sent to Esper Engine
	Events Offered	Number of events processed by the ESA service since the last service start
	Offered Rate	The rate that the ESA service processes current events / The maximum rate that the ESA service processed events.
	Status	Shows the status of the deployment. A status of Active means that the deployment is active. A status of Inactive means that there was probably an error starting the deployment. Check the error log file for more information: <code>/var/log/netwitness/correlation-server/correlation-server.log</code> .
Rule Stats	Rules Enabled	Number of rules enabled
	Rules Disabled	Number of rules disabled
	Events Matched	Total number of events matched to all rules on the ESA service
Alert Stats	Notifications	The total number of notifications sent by email, SNMP, syslog, or script for the deployment. (ESA SNMP notifications are not supported in NetWitness Platform version 11.3 and later.)
	Message Bus	The total number of alerts sent to Respond for the deployment

Deployed Rule Stats Panel

The Deployed Rule Stats panel provides details on the rules that are deployed on the ESA service. The following figure shows the Deployed Rule Stats panel.



Deployed Rule Stats

● Enable ○ Disable See [Health & Wellness](#) to monitor overall memory usage.

<input type="checkbox"/>	Enabled	Name ^	Rule Type	Trial Rule	Last Detected	Events Matched	Memory Usage	CPU %
<input checked="" type="checkbox"/>	●	AWS Critical VM Modified	Esper	Yes		0	0 bytes	0.017
<input type="checkbox"/>	●	AWS Permissions Modified Followed By Instan...	Esper	Yes		0	168 bytes	0.069
<input type="checkbox"/>	●	Account Added to Administrators Group and ...	Esper	Yes		0	168 bytes	0.134
<input type="checkbox"/>	●	Account Removals From Protected Groups on ...	Esper	Yes		0	64 bytes	0.047
<input type="checkbox"/>	●	Aggressive Internal Database Scan	Esper	Yes		0	64 bytes	0.087
<input type="checkbox"/>	●	Aggressive Internal NetBIOS Scan	Esper	Yes		0	4.22 KB	0.851
<input type="checkbox"/>	●	Aggressive Internal Web Portal Scan	Esper	Yes		0	1016 bytes	0.253
<input type="checkbox"/>	●	BYOD Mobile Web Agent Detected	Esper	Yes		0	64 bytes	0.079
<input type="checkbox"/>	●	Backdoor Activity Detected	Esper	Yes		0	0 bytes	0.014
<input type="checkbox"/>	●	Cerber Ransomware	Esper	Yes		0	896 bytes	0.314

« < | Page 1 of 1 | > » | Page Size 100 | 1 Selected Displaying 1

The table lists the various parameters in the view and their description.

Parameters	Description
<input checked="" type="radio"/> Enable	Enables a rule that was disabled.
<input type="radio"/> Disable	Disables a rule that was enabled.
Health & Wellness link	Enables you to monitor overall memory usage and health of your ESA Correlation service.
Enabled	<p>Indicates whether the rule is enabled or disabled.</p> <p>A green circle icon <input checked="" type="radio"/> indicates that the rule is enabled.</p> <p>A white circle icon <input type="radio"/> indicates that the rule is disabled.</p> <p>If a disabled rule has an error message, it shows  in the Enabled field. Hover over the icon to view the error message tooltip. The following example shows that the rule was disabled because it exceeded the configured memory threshold for that rule.</p> 
Name	Name of the ESA rule.
Rule Type	(This field applies to version 11.3 and later.) Endpoint indicates a rule from the Endpoint Risk Scoring Bundle and Esper indicates Esper-specific rules, such as Rule Builder and Advanced EPL rules.
Trial Rule	Indicates if the rule is running in trial rule mode.
Last Detected	The last time alert was triggered for the rule.
Events Matched	The total number of events that matched the rule.
Memory Usage	The total amount of memory used by the rule. <p>Note: The Endpoint Risk Scoring Rules Bundle rules do not show memory usage.</p>
CPU %	The percentage of the deployment CPU used by the rule. For example, a deployment with 1 rule shows 100% CPU usage for that rule and a deployment with two equally CPU heavy rules show 50% each. (This field is available in version 11.5 and later.) <p>Note: The Endpoint Risk Scoring Rules Bundle rules do not show CPU usage.</p>

Settings Tab

This topic describes the components of the  (Configure) > ESA Rules > Settings tab. In the Settings tab, you can perform the following tasks:

- MetaKey Referances
- Viewing the List of Meta Entities
- Enabling Meta Entity in the ESA Correlation Server
- Building Rules with Custom Meta Entities
- Configure a data enrichment source

What do you want to do?

Role	I want to ...	Show me how
Content Expert	Configure an in-memory table as an enrichment source. (Recurring In-Memory Tables are no longer supported in version 11.3 and later.)	Configure an In-Memory Table as an Enrichment Source
Content Expert	Configure a Context Hub list as an enrichment source.	Configure a Context Hub List as an Enrichment Source

Related Topics

- [Add a Data Enrichment Source](#)
- [Rule Builder Tab.](#)

Quick Look

The following figure shows the Meta Key References section in the Settings tab.

Name ^	Type
OS	string[]
access_point	string
accesses	string
action	string[]
ad_computer_dst	string
ad_computer_src	string
ad_domain_dst	string
ad_domain_src	string
ad_username_dst	string
ad_username_src	string
agent_id	string
alert	string[]
alert_id	string[]
alias_host	string[]
alias_ip	string[]
alias_ipv6	string[]

Meta Key References

The Meta Key References section lists each Meta key used by the ESA Correlation server and the type of value the key requires.

Meta key entities are configured to be a part of Event Schema. After you upgrade from RSA NetWitness® Platform 11.5, or previous releases to RSA NetWitness® Platform 11.6, you can enable the string [] meta keys entities. You can create rules and configure alerts based on the meta key entities that you select. You can also add meta entities to create rules. The meta entities retrieves data from the data sources to trigger alerts.

Note: Only the meta keys comprising of string [] are supported.

Viewing the List of Meta Entities

Perform the following steps to view the list of meta entities that are available in the system:

1. Click to **Configure** > **ESA RULES** > **Settings** > **Meta Key References**.


The Meta Key References page displays all the meta keys that are enabled.

2. Enter the name of the meta key in the Search text box to view specific meta entities.


Enabling Meta Entity in the ESA Correlation Server

After an upgrade to NetWitness 11.6 or later, you need to enable the meta entities in the ESA Correlation server to create rules and configure alerts based on the meta key entities.

To enable meta entities in the ESA Correlation server:

1. Click  (Admin) **Admin** > **Services**.
2. In the Services view, select the ESA Correlation service.
3. Click the settings icon and go to **View** > **Explore**.

The ESA Correlation Configuration page lists all the services.



Path	Value	Service
/rsa/correlation/stream		esaprimary - ESA Correlation
dots-to-underscores	true	
enable-meta-entity	true	
event-batch-size	1000	
event-enrichment-queue-size	10	
event-enrichment-thread-pool-size	8	
event-polling-timeout-in-milli-seconds	1000	
event-source-id	true	
filter		
idle-retry-interval	10	
lag-time	15 MINUTES	
lowercase		
max-sessions	10000	
mechanism	AGGREGATION	
meta-entities	checksum.all , compromise.all , domain.all , host.all , ip.all , outcome.all , user.all	

4. Click **enable-meta-entity** and set it to **true**.

Note: By default, **enable-meta-entity** is set to **false**. Ensure that you enable only the specific meta entity using which you want to build a rule. Enabling multiple meta entities affects the system performance.

5. Add custom meta entities by appending or editing the meta-entities field values under explore view.

In the Explore view node list for the ESA Correlation service, select **Correlation** > **Stream**. All the stream settings are listed. You can add custom meta entities to the list as per the requirement. Once you enable custom meta entities, you can view them under Meta Key References. Once you enable the custom meta key entities, the rules will start evaluating the meta key entities and start to evaluate and trigger alerts as per the configured conditions. For more information, see [Viewing the List of Meta Entities](#)

The following meta entities are listed in the ESA Correlation Configuration page:

- checksum.all
- compromise.all
- domain.all
- host.all
- ip.all
- outcome.all
- user.all

Meta Entity Usage Example

ip_all, and *host_all* are some of the examples of meta entities.

For example, the *ip_all* meta entity can be used to check for any malicious IPs. It combines all the IPv4 meta keys (*ip_src*, *alias_ip*, *ip_dst*, and *forward_ip*) and creates a string array to retrieve information and trigger alerts.

Note: You can add and enable custom meta entities (string arrays) to the list.

Building Rules with Custom Meta Entities

This topic provides instructions to define rule criteria in Rule Builder by adding statements. A statement is a logical grouping of rule criteria in the Rule Builder. You add statements to define what a rule detects.

The following graphic shows an example of a Rule Builder statement using meta entities.

Every statement contains a key and value. Then, you build logic around the pair by selecting an option in each other field.

Build a Statement

Define a rule condition by adding one or more statements. For each statement, define the keys, operators, and values that will trigger the rule. If the contents of the value field include more than one value, you must specify that it should be evaluated as an array.

Name *

Select any or all...

	Key	Operator	Value	Ignore Case?	Array?
<input checked="" type="checkbox"/>	ip_all	Select Operat...	Enter Value	<input type="checkbox"/>	<input type="checkbox"/>

Prerequisites

To build a rule statement, you must know the meta key and the meta value. For a complete list of meta keys, go to **Configure > ESA Rules > Settings > Meta Key References**. For more information about building rules with meta entities, see [Rule Builder Tab](#).

Enrichment Sources

In the Enrichment Sources section, you can use the following external data sources:

- GeoIP
- In-Memory Table (Add hoc only - Recurring In-Memory Tables are no longer supported in version 11.3 and later.)
- Context Hub

The following figure shows the Enrichment Sources section in the Settings tab.

The screenshot displays the RSA NetWitness Platform interface. The top navigation bar includes 'Investigate', 'Respond', 'Users', 'Hosts', 'Files', 'Dashboard', and 'Reports'. Below this, a secondary navigation bar lists 'LIVE CONTENT', 'SUBSCRIPTIONS', 'INCIDENT RULES', 'INCIDENT NOTIFICATIONS', 'ESA RULES', 'CUSTOM FEEDS', 'EVENT RULES', and 'LOG PARSER RULES'. The 'Settings' tab is selected, and the 'Enrichment Sources' section is active. The interface shows a table with the following data:

Enabled	Name	Type	Description	Last Modified	Actions
<input type="checkbox"/>	Default GeoIP	GeoIP	Default Geo IP Enrichment Source. This cannot be edited.	2020-06-02 04:57:36	

The interface also includes a search bar, a pagination control showing 'Page 1 of 1', and a page size dropdown set to '100'. The bottom of the screen displays 'RSA NETWITNESS PLATFORM' and the version '11.5.0.0'.