



ESA Configuration Guide

for Version 11.2



Copyright © 1994-2018 Dell Inc. or its subsidiaries. All Rights Reserved.

Contact Information

RSA Link at <https://community.rsa.com> contains a knowledgebase that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

For a list of RSA trademarks, go to www.emc.com/legal/emc-corporation-trademarks.htm#rsa.

License Agreement

This software and the associated documentation are proprietary and confidential to Dell, 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 Dell.

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 RSA Link. 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

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

August 2018

Contents

Event Stream Analysis Overview	5
Configure ESA Correlation Rules	7
Prerequisites	7
Procedure	7
Result	7
Step 1. Add a Data Source to an ESA Service	7
Prerequisites	8
Procedures	8
Step 2. Configure Advanced Settings for an ESA Service	10
Procedures	10
Configure ESA Analytics	12
Configure the Whois Lookup Service	12
Prerequisites	12
Mapping ESA Data Sources to Analytics Modules	15
Module Deployment Example - Two ESAs	15
Module Deployment Example - One ESA	16
Prerequisites	17
Create ESA Analytics Mappings	18
Deploy ESA Analytics Mappings	21
Update a Mapping	22
Undeploy a Mapping	22
Delete a Mapping	22
Change the Warm-up Period and Lag Time	23
Additional ESA Correlation Rules Procedures	25
Change Memory Threshold for Trial Rules	25
Prerequisites	25
Configure ESA to Use a Memory Pool	26
Result	29
Configure ESA to Use Capture Time Ordering	29
Capture Time Order Workflow	30

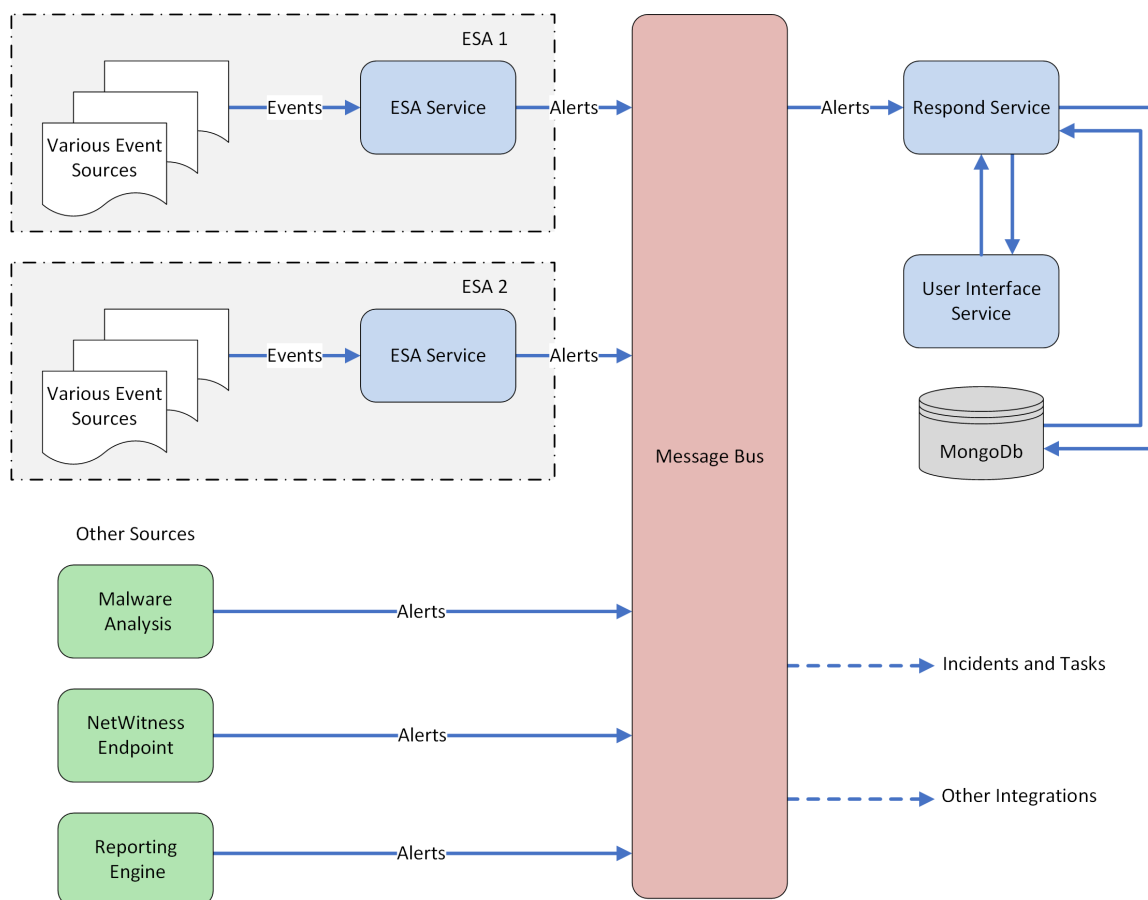
Prerequisites	31
Procedures	31
Troubleshooting Tips	33
Start, Stop, or Restart ESA Service	34
Start ESA Service	34
Stop ESA Service	34
Restart ESA Service	34
Audit Logs and Verify ESA Component Versions and Status	34
Audit Log Rules	34
Verify ESA Server Version	36
References	37
Services Config View Data Sources Tab	38
Workflow	38
What do you want to do?	39
Related Topics	39
Quick Look	39
Services Config View Advanced Tab	42
Workflow	42
What do you want to do?	43
Related Topics	43
Quick Look	43
Whois Lookup Service Configuration	46
What do you want to do?	46
Related Topics	46
Quick Look	47
ESA Analytics Mappings	50
Workflow	50
What do you want to do?	51
Related Topics	51
Quick Look	51
Module Settings	56
What do you want to do?	56
Related Topics	56
Quick Look	56

Event Stream Analysis Overview

RSA NetWitness® Platform Event Stream Analysis (ESA) 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.

ESA's advanced Event Processing Language allows you to express filtering, aggregation, joins, pattern recognition, and correlation across multiple disparate event streams. Event Stream Analysis helps perform powerful incident detection and alerting.

The following diagram shows the high-level data workflow:



There are two ESA services that can run on an ESA host:

- Event Stream Analysis (ESA Correlation Rules)
- Event Stream Analytics Server (ESA Analytics)

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

The second service is the ESA Analytics service, which is used for Automated Threat Detection. Because the ESA Analytics service uses preconfigured ESA Analytics modules for Automated Threat Detection, you do not have to create or download rules to use Automated Threat Detection.

ESA Analytics services use query-based aggregation (QBA) to collect filtered events for the ESA Analytics modules from Concentrators. Only the data required by a module is transferred between the Concentrator and the ESA Analytics system. For example, using a Suspicious Domains ESA Analytics module, such as C2 for Packets ([http-packet](#)), an ESA Analytics service can examine your HTTP traffic to determine the probability that malicious activity is occurring in your environment.

Configure ESA Correlation Rules

This topic provides high-level tasks to configure RSA NetWitness Platform Event Stream Analysis (ESA) Correlation Rules using the Event Stream Analysis service.

Prerequisites

Make sure that you:

- Install the Event Stream Analysis service in your network environment.
- Install and configure one or more Concentrators in your network environment.

Procedure

The following table shows the high level tasks required to configure ESA Correlation Rules.

Tasks	Reference
1. Add a Concentrator as data source to the Event Stream Analysis service.	Refer to Step 1. Add a Data Source to an ESA Service .
2. Configure notifications for the Event Stream Analysis service.	Refer to "Notification Methods" in the <i>Alerting with ESA Correlation Rules User Guide</i> .
3. Download Event Stream Analysis content using Live.	Refer to "Download Configurable RSA Live Rules" in the <i>Alerting with ESA Correlation Rules User Guide</i> .
4. (Optional) Advanced configuration for the Event Stream Analysis service.	Refer to Step 2. Configure Advanced Settings for an ESA Service .

Result

The Event Stream Analysis service is configured and you can now add ESA Rules for event processing and alerting. For information on adding ESA Rules, see "Add Rules to the Rule Library" in the *Alerting with ESA Correlation Rules User Guide*.

Step 1. Add a Data Source to an ESA Service

This topic describes how to add a new or existing data source to the Event Stream Analysis service.

An ESA service ingests data from a Concentrator to detect incidents and alert the user. For ESA to analyze data, you need to configure the sources from which the ESA will read data. Use the procedures in this topic to add data sources for your ESA.



Prerequisites

You must have one or more Concentrators configured in NetWitness Platform.

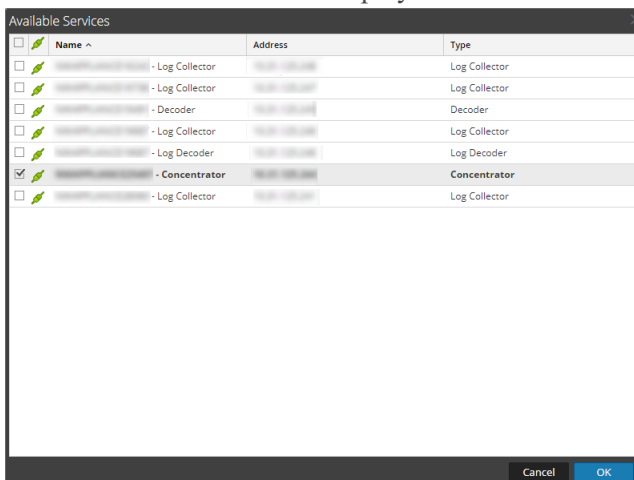
The Event Steam Anaysis service must be installed and running on NetWitness Platform.

Procedures

Add an Existing Service as Data Source

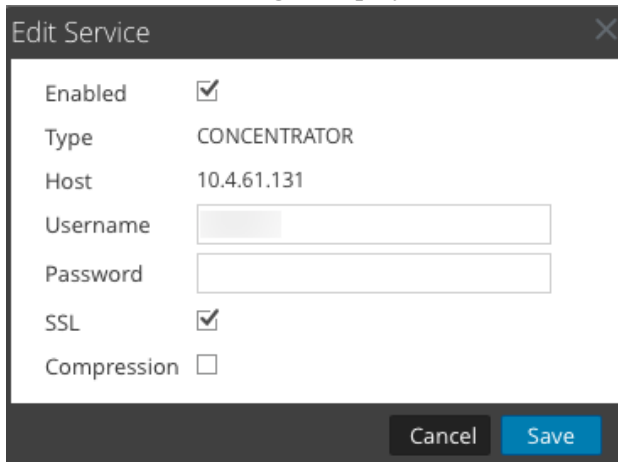
1. Go to **ADMIN > Services**.
The Services view is displayed.
2. In Services view, select an ESA service and select  > **View > Config**.
3. On the **Data Sources** tab, click  .

The available services are displayed as shown in the following figure.



4. Select a Concentrator service and click **OK**.

The Edit Service dialog is displayed.



Enabled	<input checked="" type="checkbox"/>
Type	CONCENTRATOR
Host	10.4.61.131
Username	<input type="text"/>
Password	<input type="password"/>
SSL	<input checked="" type="checkbox"/>
Compression	<input type="checkbox"/>


Cancel Save

5. Click **Enable** to enable (or disable) the data source (it is enabled by default when adding a new service).
6. Enter a valid **Username** and **Password** for the service.
7. Click to enable or disable the **SSL** or **Compression** options.
8. Click **Save** to save the configuration and close the Edit Service dialog.
9. Click **Apply** to complete the change on the **Data Sources** tab.
The service is added to the list of services in the **Data Sources** tab.

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

Edit Settings for a Data Source

To edit settings, including the username and password, for a configured data source:

1. Go to **ADMIN > Services**.
The Services view is displayed.
2. In the **Services** view, select a Concentrator service.
3. Click .
The Edit Service dialog is displayed (see previous figure).
4. Modify the settings as desired, including entering a new username and password. The username field will be prepopulated with the currently configured username. To change the password, enter a new password in the password field. If you leave the password field blank, the previously configured password will continue to be used.
5. Click **Save** to save the changes and close the Edit Service dialog.
6. Click **Apply** to complete the change on the **Data Sources** tab.

Step 2. Configure Advanced Settings for an ESA Service



This topic provides instructions to configure advanced settings for an Event Stream Analysis service.


In the Advanced view, you can configure advanced settings to improve performance, to preserve events for rules with multiple events, to buffer events in memory, and to specify the number of events to be stored on the ESA.

Procedures

Configure Advanced Settings

To access the Advanced view and configure advanced settings for an ESA service:

1. Go to **ADMIN > Services**.
The Services view is displayed.
2. In Services view, select an ESA service and   > **View > Config**.
3. Select the **Advanced** tab.
The Advanced view is displayed.



Configure Alert Engine Settings

In the Alert Engine section, you specify values to preserve events for rules that choose multiple events. The following figure shows the Alert Engine section.

Alert Engine

Max Constituent Events

Forward Alerts On Message Bus

Debug Rules?

Apply

To configure Alert Engine settings:

1. In the Alert Engine section, specify a value for **Max Constituent Events**. The default value is 100.
2. Select **Debug Rules?** to enable debugging rules.
3. If you want alerts to be sent to Message Bus and Respond, select the **Forward Alerts On Message Bus** option.
4. Click **Apply** to save the changes and put them into effect immediately.

Note: For more information see [Services Config View Advanced Tab](#).

Configure Event Stream Engine Settings

In the Event Stream Engine section, you specify details to improve performance. The following figure shows the Event Stream Engine section.

Event Stream Engine

Max Pattern Subexpressions

Apply

To configure Event Stream Engine settings:

1. In the Event Stream Engine section, specify **Max Pattern Subexpressions**.
2. Click **Apply** to save the changes and put them into effect immediately.

Note: For more information see [Services Config View Advanced Tab](#).

Configure ESA Analytics

This section provides high-level tasks to configure ESA Analytics services for RSA NetWitness® Platform Automated Threat Detection. The Automated Threat Detection functionality enables you to analyze the data that resides on one or more Concentrators by using preconfigured ESA Analytics modules, such as Suspicious Domains. For example, using a Suspicious Domains module, an ESA Analytics service can examine your HTTP traffic to determine the probability that malicious activity is occurring in your environment.

There are two ESA services that can run on an ESA host:

- Event Stream Analysis (ESA Correlation Rules)
- Event Stream Analytics Server (ESA Analytics)

The first service is the Event Stream Analysis service that creates alerts from ESA rules, also known as ESA Correlation Rules, which you create manually or download from Live. The second service is the ESA Analytics service, which is used for Automated Threat Detection and is configured in this section. Because the ESA Analytics service uses preconfigured ESA Analytics modules for Automated Threat Detection, you do not have to create or download rules to use it.

There are currently two ESA Analytics modules available and they are both for Suspicious Domains:

- C2 for Packets (http-packet)
- C2 for Logs (http-log)

Configure the Whois Lookup Service

The RSA NetWitness Platform Automated Threat Detection functionality enables you to automatically analyze data sources by using preconfigured ESA Analytics modules. An ESA Analytics module is a pipeline composed of activity objects that enrich an event with additional information through mathematical computations. ESA Analytics services process these modules to identify advanced threats.

The Whois Lookup service configuration is required for the Suspicious Domains modules.

Note: (Important) RSA strongly recommends that you configure the Whois Lookup service for accuracy in Automated Threat Detection scoring.

Prerequisites

- You must have an RSA Live account to use the Whois Lookup service.
- The ESA Analytics Server service must be available (shows a green circle) in the ADMIN > Services view.

If you configured a Live account in the Live Services panel (ADMIN > System > Live Services), the Whois Lookup Service is automatically configured for you. You only need to check the connection of the Whois Lookup service.

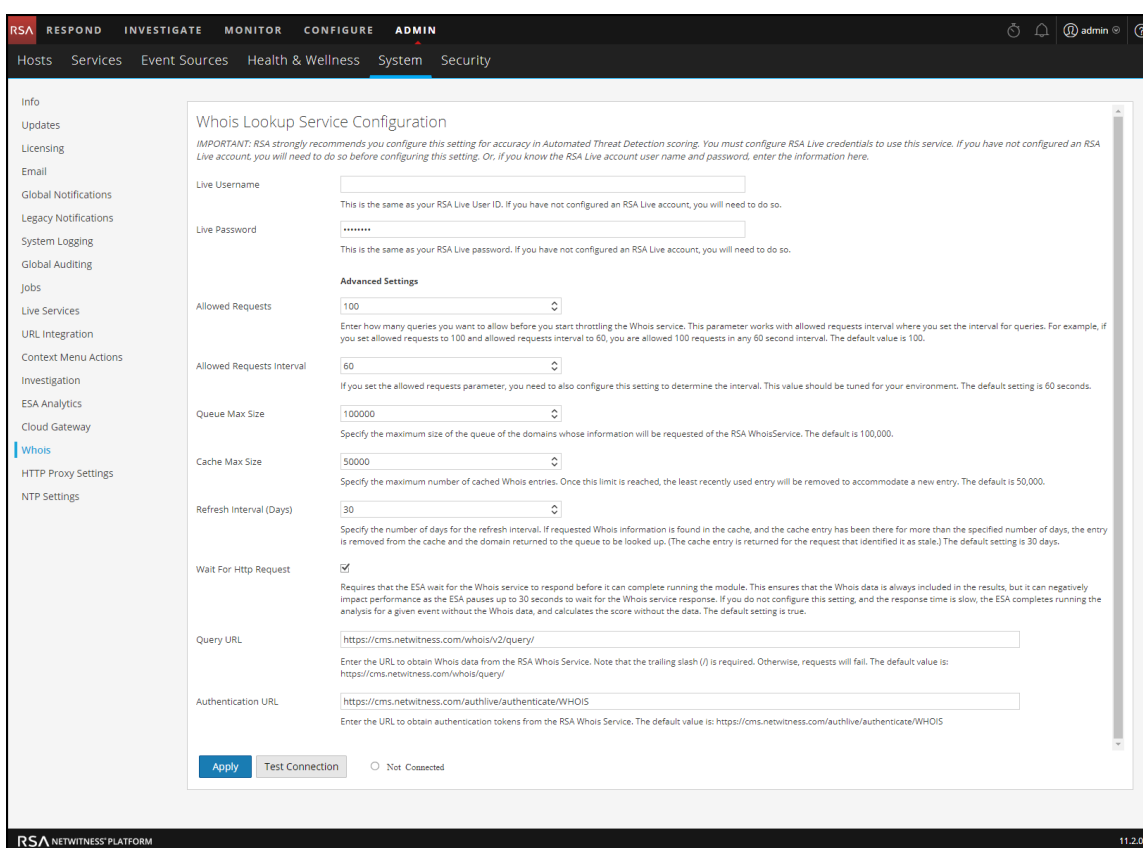
Note: If you do not have an RSA Live account, you can create one at the RSA Live Registration Portal:

<https://cms.netwitness.com/registration/>

The *Live Services Management Guide* provides additional information.

To configure the Whois Lookup service:

1. Go to **ADMIN > System**.
2. In the options panel, select **Whois**.
3. In the **Whois Lookup Service Configuration** panel, check to see if the Whois Lookup service is connected. At the bottom of the panel, a connected service shows a green circle next to **Connected**:



If it is connected, you are finished with the configuration and you can skip the remaining steps. To adjust the advanced settings, go to step 5.

If the service is not connected, continue to step 4.

4. In the **Live Username** and **Live Password** fields, enter your RSA Live account credentials to access the RSA Whois server.
5. If necessary, you can adjust the advanced settings. However, RSA recommends that you use the default values. [Whois Lookup Service Configuration](#) provides additional details.

6. To test your connection, click **Test Connection**.

A successful connection shows a green circle next to **Connected**:



7. Click **Apply** to save your changes.

Mapping ESA Data Sources to Analytics Modules

This topic tells Administrators how to map specific ESA Analytics modules to multiple data sources and ESA Analytics services, which can make processing more efficient.

You can analyze the data that resides on one or more Concentrators with the RSA NetWitness Platform Automated Threat Detection functionality by selecting a preconfigured ESA Analytics module. The data analyzed by these modules is used to identify advanced threats. To better utilize your network resources and reduce unnecessary data flow, you can map multiple data sources, such as Concentrators, to multiple ESA Analytics services in order to process data more efficiently and take advantage of additional capacity.

An *ESA Analytics module* is a pipeline composed of activity objects that enrich an event with additional information through mathematical computations. ESA Analytics modules reside within ESA Analytics services.

When you deploy your mapping, the selected ESA Analytics services use query-based aggregation to collect the appropriate filtered events for the selected module from the Concentrators. Query-based aggregation is a predefined query that only transfers data for the selected ESA Analytics module. Only the data required by the module is transferred between the Concentrator and the ESA Analytics system.

There are currently two ESA Analytics modules available for Suspicious Domains: C2 for Packets ([http-packet](#)) and C2 for Logs ([http-log](#)).

Module Deployment Example - Two ESAs


To take advantage of your additional Concentrator capacity, you can map an ESA Analytics module to an ESA Analytics service and deploy it to analyze data from multiple data sources at the same time.

For example, if you have three Concentrators and two ESA Analytics services, you can create and deploy the following mappings:

- Map Module 1 to the Concentrator 1 and 2 sources and the ESA Analytics 1 service. ESA Analytics Service 1 analyzes Module 1 filtered events from Concentrators 1 and 2.
- Map Module 2 to the Concentrator 2 and 3 sources and the ESA Analytics 2 service. ESA Analytics Service 2 processes Module 2 filtered events from Concentrators 2 and 3.

In this example, Module 1 represents an ESA Analytics module, such as C2 for Packets ([http-packet](#)) and Module 2 represents another ESA Analytics module, such as C2 for Logs ([http-logs](#)) in another location.

Module Deployment Example – Two ESAs



This example shows how both services can process data from the same Concentrator. Notice that ESA Analytics Services 1 and 2 can both process data from Concentrator 2. ESA Analytics Service 1 queries data for Module 1 events and ESA Analytics Service 2 queries different data for Module 2 events.


Module Deployment Example - One ESA

In addition to creating module mappings that are processed by different ESA Analytics services, you can map more than one module to the same ESA Analytics service.

For example, if you have three Concentrators and one ESA Analytics service, you can create and deploy the following mappings:

- Map Module 1 to the Concentrator 1 and 2 sources and the ESA Analytics 1 service. ESA Analytics Service 1 analyzes Module 1 filtered events from Concentrators 1 and 2.
- Map Module 2 to the Concentrator 2 and 3 sources and the ESA Analytics 1 service. ESA Analytics Service 1 also processes Module 2 filtered events from Concentrators 2 and 3.

Module Deployment Example – One ESA



This example shows how one service can process data from more than one module. Notice that ESA Analytics Service 1 can process data from Concentrators 1 and 2 for Module 1. It also processes data from Concentrators 2 and 3 for Module 2. ESA Analytics Service 1 queries data for Module 1 events and queries different data for Module 2 events.

Caution: Ensure that all NetWitness Platform host services are in sync with a consistent time source.

Prerequisites

- All NetWitness Platform host services must be in sync with a consistent time source.
- The Concentrator hosts and services must be discovered and available in the NetWitness Platform user interface.
- All module-specific requirements must be followed.
 - For Suspicious Domains:
 - Configure log settings (Suspicious Domains for Logs only)
 - Create a whitelist using the Context Hub service.
 - [Configure the Whois Lookup Service](#).
 - Verify that the C2 incident rule is enabled and monitor it for activity.
 - Verify that the incidents are grouped by Suspected C&C.

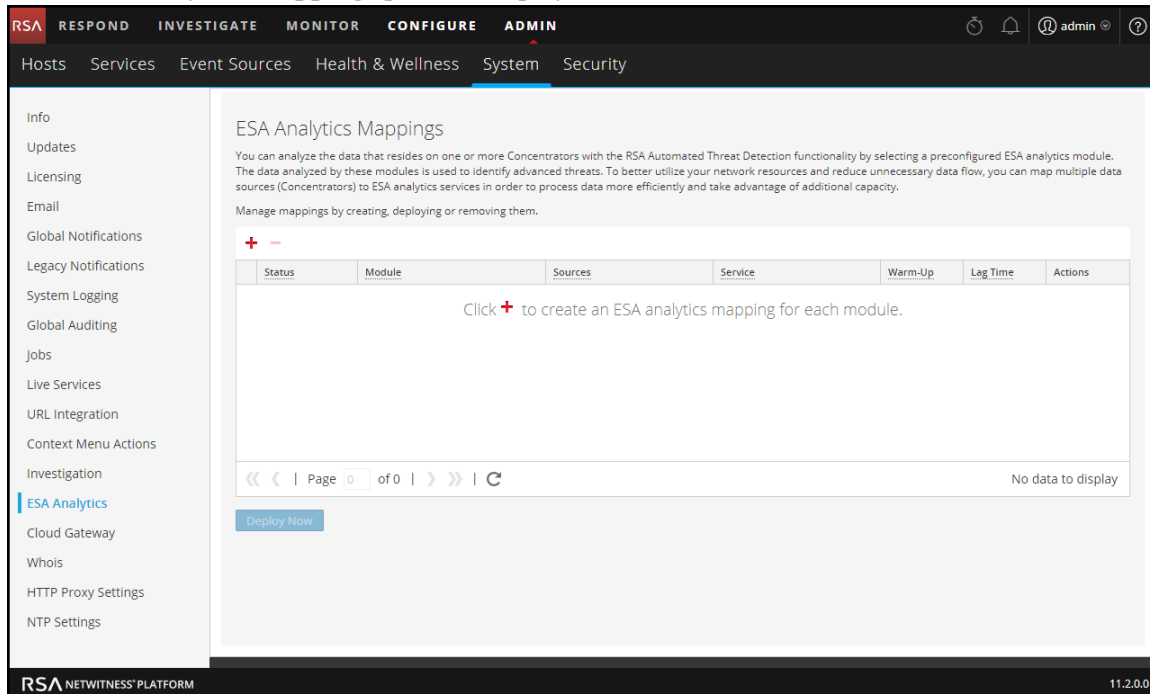
For step-by-step procedures, see the *NetWitness Platform Automated Threat Detection Configuration Guide*.

Create ESA Analytics Mappings

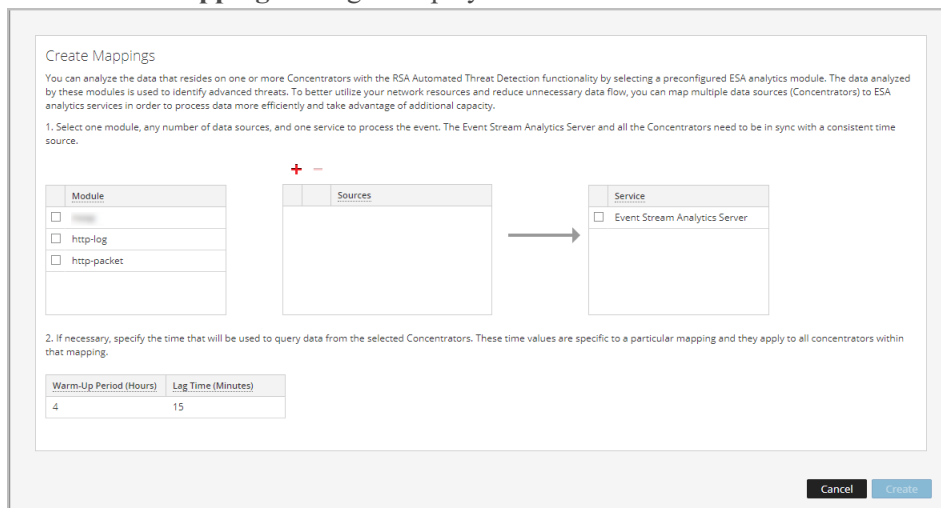
The following procedure tells you how to map ESA Analytics modules to sources and services. After creating and reviewing the mappings, you deploy them so that they can start aggregating data.

1. Go to **ADMIN > System**, and in the options panel, select **ESA Analytics**.

The **ESA Analytics Mappings** panel is displayed.



2. Click **+** to create an ESA Analytics mapping. Create a separate mapping for each module. The **Create Mappings** dialog is displayed.




3. In the **Module** list, select a module.

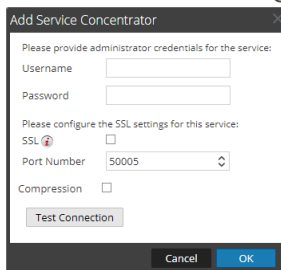
4. Configure one or more data sources (Concentrators) for your mappings. Do the following for each Concentrator:

- a. Click **+**.


The Available Services dialog shows the data sources that are available from the Admin > Services view.



- b. In the **Available Services** dialog, select a Concentrator and click **OK**. The **Add Service** dialog is displayed.



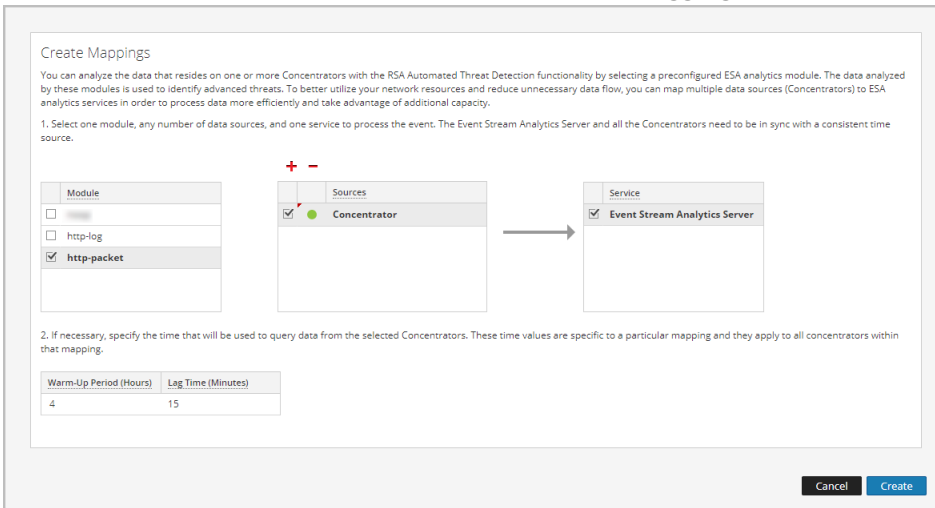
- c. In the **Add Service** dialog, type the Administrator username and password for the Concentrator.
- d. Click **Test Connection** to make sure that it can communicate with the ESA Analytics service.



- e. Click **OK**.

After you configure your data sources and they appear in the Sources list, you can reuse them for additional mappings.

- In the **Sources** list, select one or more data sources to aggregate the data for the module.



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


- In the **Service** list, select an ESA Analytics service to process the data for the module.
- If necessary, specify the time that will be used to query data from the selected Concentrators:

Field	Description
Warm-up Period (Hours)	<p>Specifies a warm-up duration (in hours). A warm-up period is required to allow Automated Threat Detection to "learn" your traffic. The warm-up period should run when typical traffic is running. During this time, alerting for your module mapping is suppressed. The Warm-up Period primes the module with historical data and guarantees that the specified number of hours of data collection completes before sending alerts.</p> <p>RSA provides preconfigured ESA Analytics modules. Each module type has a default warm-up period defined, which you can adjust to your environment, if necessary. After this warm-up period, alerts can be viewed.</p> <p>For more information about Warm-up Period and Lag time, see Module Settings.</p>

Field	Description
Lag Time (Minutes)	<p>Specifies a constant time delay in minutes, which is added to avoid losing events being processed by the data sources during periods of heavy activity. For example, Concentrator performance varies depending on factors such as incoming load, ongoing queries, and indexing. Due to these factors, a Concentrator may not aggregate events in real-time, which leads to the delay.</p> <p>The Lag parameter gives the Concentrator a chance to finish aggregating all of the data.</p> <p>After the warm-up period completes, data aggregation continues at Current (System) Time - Lag Time. This is useful when a Concentrator is slow in aggregating data. The Lag time guarantees that the module does not process data that arrives to the Concentrator within the Lag time window so there is adequate delay to ensure all events that get generated in the enterprise can be processed by the module.</p> <p>For example, if Lag time is 30 minutes, and the current time is 2:00 PM, the Concentrator starts pulling records at 1:30 PM. The Lag time window, 30 minutes in this example, remains constant as time advances. When the current time advances to 2:01 PM, the Concentrator pulls the next minute of data at 1:31 PM, and so on.</p> <p>Important: The Lag time defines the buffer between the current time and the time when the module ingests the data.</p> <div style="border: 1px solid yellow; padding: 5px;"> <p>Caution: RSA recommends that Administrators adjust the Lag parameter dynamically based on the performance of each of the individual Concentrators to avoid missing any events during aggregation.</p> </div> <p>For more information about Warm-up Period and Lag time, see Module Settings.</p>

8. Click **Create**.

The mappings that you create appear in the list of existing mappings with a status of **Undeployed**.



Important: To start a module so that it starts aggregating data, you need to deploy it.

Deploy ESA Analytics Mappings

After you create your mappings, you need to deploy them in order to start aggregating data for the modules.

1. In the list of mappings, verify that the status of the mappings that you want to deploy show as **Undeployed**.
2. Select one or more mappings with a status of Undeployed and select **Deploy Now**.
All selected mappings in the Undeployed state start to aggregate data as configured in the mapping.
The mapping status changes to **Deployed**.
You cannot deploy a mapping that has already been deployed.

Update a Mapping

You can only have one mapping per module. If you want to make changes to a deployed mapping, such as adding or removing Concentrators or changing the service, you must undeploy and delete the existing mapping and then create and deploy a new mapping for that module.

You can make the following updates to a deployed mapping without deleting it:

- Undeploy the mapping
- Change the warm-up period and lag time



You can also change the warm-up period and lag time for an undeployed module mapping.

Undeploy a Mapping

If you want to stop aggregating data for a module mapping, but you do not want to delete the mapping, you can undeploy it. This gives you the option of deploying it at a later time. When you undeploy a mapping, the specified ESA Analytics service stops pulling data from the data source for that module.

Caution: Undeploying a mapping with a status of Deployed will affect data aggregation for that module.

To undeploy a mapping:

1. In the ESA Analytics Mappings panel, select the deployed mapping that you want to undeploy.
2. In the **Actions** column, select   > **Undeploy**.
The status changes from Deployed to Undeployed and data aggregation stops.

Delete a Mapping

You can delete a mapping with a status of Undeployed at any time. Since a mapping in the Undeployed state is not running, it does not affect data aggregation.

You should undeploy a mapping with a status of Deployed before deleting it. Undeploying and deleting a mapping clears the configuration on the ESA server, reverts the deployment for that mapping, and stops pulling data from the data source for that module.

Caution: Undeploying and deleting a mapping will affect data aggregation for that module.

To delete a mapping:


1. In the ESA Analytics Mappings panel, select the mapping that you want to delete. You can only delete one mapping at a time.

2. Click  .

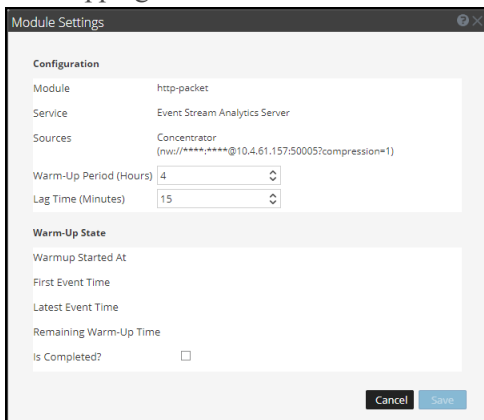
Change the Warm-up Period and Lag Time

You may want to adjust the warm-up period for a specific module mapping. For example, after the warm up period is complete, you can increase the warm-up period setting to allow additional warm-up time. You can even increase the warm-up period when your module mapping is actively warming up.

If necessary, you can change the lag time for the module. The lag time defines the buffer between the current (system) time and the time when the module ingests the data.

1. In the ESA Analytics Mappings panel, select the mapping that you want to change and in the **Actions** column, select  > **Edit Module**.




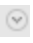
The Module Settings dialog shows the selected module, ESA Analytics service, and data sources for the mapping. The data sources show the URLs used to communicate with ESA.



2. Review the **Warm-Up State** section to determine the current warm-up state:
 - **Warm Up Started At** - The time when the first event was processed by the ESA Analytics module from the data source.
 - **First Event Time** - The time that the first event occurred. The warm-up time is based on this time.
 - **Latest Event Time** - The time that the latest event occurred.
 - **Remaining Warm Up Time** - The number of hours remaining in the warm-up period.
 - **Is Completed?** - Indicates whether the warm-up period is complete. If it is true, the warm-up period is complete. If it is false, the module is still warming up and you can view the number of hours remaining in the Remaining Warm Up Time field.
3. In the **Configuration** section, you can update the **Warm-Up Period (Hours)** depending on whether or not the warm-up period is complete.
 - **During the warm up period** - You can add hours to the warm-up period or subtract any remaining warm-up time.
 - **The warm-up period is complete** - You can add hours to the warm-up period by adding the difference between the current time and the First Event Time to the hours that you want to add. For example, a warm-up period of 10 hours is complete and the First Event Time shows 12:00:00.

The current time is 16:00:00 (4 hours later) and you want to add 5 more hours to the warm-up time. To do this, you need to add 9 hours ($4+5=9$) to the warm-up period of 10, so you would set the new warm-up period to 19 hours.

You cannot decrease the warm-up period if it is complete, unless you delete the mapping and create a new one.

4. If necessary, you can adjust the **Lag Time (Minutes)** to give the Concentrators in the mapping additional time to finish aggregating all of the data.
5. Click **Save**.
Changes **DO NOT** take effect immediately. For the settings to take effect, you need to undeploy and re-deploy the mapping.
6. To undeploy the mapping, in the ESA Analytics Mappings panel, select the mapping that you want to undeploy and then select   > **Undeploy**.
Data aggregation stops for the selected mapping.
7. To re-deploy the mapping, select the mapping that you want to deploy and then select   > **Deploy**.
The selected mapping deploys and starts to aggregate data as configured in the mapping.

Additional ESA Correlation Rules Procedures

This topic is a collection of individual procedures, which an Administrator may perform at any time and they are not required to complete the initial setup of ESA Correlation Rules.

Use this section when you are looking for instructions to perform a specific task after the initial setup of ESA.

- [Change Memory Threshold for Trial Rules](#)
- [Configure ESA to Use a Memory Pool](#)
- [Configure ESA to Use Capture Time Ordering](#)
- [Start, Stop, or Restart ESA Service](#)
- [Audit Logs and Verify ESA Component Versions and Status](#)

Change Memory Threshold for Trial Rules

This procedure is optional and applies only to ESA Correlation Rules.

Administrators can increase or decrease the memory threshold for trial rules. Threshold refers to the ESA memory usage, which includes ESA base memory, trial rules, and non-trial rules. When the threshold is exceeded, all deployed trial rules on an ESA service are disabled.

You use trial rules to see if a rule runs efficiently and does not use excessive memory, which can impact performance or force the service to shut down.

By default, the memory threshold is 85, which is the percentage of Java Virtual Memory (JVM).



- The memory threshold is per ESA, not per rule.
- When the memory threshold is exceeded, all trial rules running on the ESA are automatically disabled.
- The ESA configuration has two parameters for trial rules:
 - `MemoryThresholdforTrialRules`
 - `MemoryCheckPeriod`, which has a default value of 300 seconds

For more information, see "Work with Trial Rules" in the *Alerting with ESA Correlation Rules User Guide*.


Prerequisites

A role with administrative privileges must be assigned to you.

To change memory threshold for trial rules:

1. Log on to NetWitness Platform as admin.
2. Go to **ADMIN > Services**.
3. Select the ESA service and select   > **View > Explore**.

- On the left, select **CEP > Module > configuration**.



- In the right panel, in **MemoryThresholdForTrialRules** type a percentage of JVM that trial rules on the ESA can not exceed.
The new memory threshold takes effect immediately.

Configure ESA to Use a Memory Pool

This procedure applies only to ESA Correlation Rules.

Administrators can configure ESA to use a memory pool. A memory pool is a customized implementation of virtual memory for events held by rules in ESA. This helps in scaling the capability of rules by an order of magnitude. When you want to create rules that cover a large time span or which are very complex, you may want to use a memory pool to handle memory more efficiently. When you use a memory pool, instead of holding all of the events in memory, they can be written to disk. This is helpful because when a rule exists that is complex or extends over a long time frame, a large number of events must be held in memor

You can configure memory pool to run in non-batch mode or batch mode:


- Non-batch mode.** In non-batch mode, events are written to disk as they enter the memory pool. To configure non-batch mode, set the **MapPoolBatchWriteSize** attribute to 1. Non-batch mode provides a more stable solution because each event is landed and fetched separately without creating memory spikes.
- Batch mode.** In batch mode, events are grouped into batches and then written to disk. To configure batch mode, set the batch size attribute **MapPoolBatchWriteSize** to a value greater than 1. Batch mode gives better performance since the disk activity for landing events to disk are optimized.

Note: Any changes to these settings will require you to restart the ESA. When ESA restarts, if any events are currently being held by the memory pool, they will be discarded upon restart.

Caution: While this feature can be very helpful in managing memory, it can impact the event processing rate of the ESA. Performance can be affected from 10 to 30 percent, depending on your rules and configuration settings.


Workflow

The following diagram shows the data flow using the memory pool for batch mode.



1. Events are added to the memory pool and references to the events are stored in the memory pool.
2. The events are then batched to be sent to disk (in non-batch mode, this step is skipped).
3. Once the batch has met the threshold, the events are written to disk (in non-batch mode no threshold is required).
4. When the EPL requires an event that was written to disk, the event is sent to the cache and used in the EPL rule.

To configure an ESA memory pool:

1. Go to **ADMIN > Services**, select your ESA service, and then  > **View > Explore**.
2. Select **CEP > EsperPool > Configuration**.
3. Enter values for the following fields:

Attribute	Description	Configuration
MapPoolPersistenceURI	Location to store the memory pool file.	<p>The default value is <code>/opt/rsa/esa/pool/esperPool</code>. RSA recommends you do not modify the default value.</p> <p>If you modify this setting to use a different partition, ensure the partition contains at least 10 times more space than the memory allocated for ESA.</p> <div data-bbox="784 501 1422 653" style="border: 1px solid yellow; padding: 5px;"> <p>Caution: If the memory pool is in use while this path is changed, an ESA restart is required. When this occurs, ESA does not discard the stored events so you must manually purge them.</p> </div>
MapPoolEnable	Enable or disable the memory pool.	<p>The default value is false. Set the value to true to enable the memory pool. Requires a restart when you enable or disable memory pool.</p>
MapPoolFlushIntervalSecs	Time interval to flush events to disk. For example, any event held in Esper longer than 15 minutes gets flushed to disk.	<p>The default value is 15 minutes. A smaller value ensures that the ESA is more stable when there are EPLs holding a large number of events in memory. A larger value (greater than 30 minutes), ensures that only relevant events required over a longer period of time are flushed to disk.</p> <div data-bbox="784 1062 1422 1241" style="border: 1px solid green; padding: 5px;"> <p>Note: Due to Java memory management design, sometimes events not held by EPL may be sent to disk. To help prevent this from occurring, you can set a higher value for <code>MapPoolFlushIntervalSecs</code>.</p> </div>
MapPoolBatchWriteSize	<p>Specify the batch size (and whether to use batch mode). The events are batched into groups and then flushed to disk.</p> <p>To use non-batch mode, set this value to 1.</p> <p>To use batch mode, set this value to greater than 1.</p>	<p>The default batch size is 100,000 events. At the end of flush interval, if the batch capacity is not reached, the batch expires in 30 seconds and all contents of the batch are written to disk as memory pool files.</p> <p>A smaller value for the batch size (for example, 10,000 events) ensures that when events are fetched from disk, they do not pose a risk of bloating the memory, which creates more stability. However, a larger batch size (100,000 events) minimizes the input/output activity when writing events to disk, which can create better performance.</p>

Attribute	Description	Configuration
MapPoolMinSize	Minimum size of the memory pool. This value is used for initialization, so it does not typically require editing.	The default value is 10,000 events. A higher value may increase performance. A lower value ensures that the system is more stable.
MapPool Persist Type	This is a view-only parameter that displays the type of optimization used.	The default value is RMSerialize .

Note: The effectiveness of this feature depends on your environment. If you write rules that require frequent access of events over a period of time, this feature may degrade performance with no or minimal improvement in scalability.

Memory pool files get deleted when all the events held in the pool file are no longer referenced by an EPL.

Result

For a simple EPL rule, ESA typically improves memory approximately 8 to 9 times.

Configure ESA to Use Capture Time Ordering

This procedure applies only to ESA Correlation Rules.

Administrators can configure the ESA to use capture time ordering when using two or more Concentrators as a source.

By default, ESA uses the ESA time stamp (time at which events are received by the ESA) to correlate events. However, ESA also supports session-ordering based on capture time (the time at which the packet or log event reached the Decoders). This feature is useful if you are correlating events from two or more Concentrators. When you have two or more Concentrators as sources, time ordering ensures that their sessions are correlated together by capture time. This ensures that sessions captured at the same time are correlated together and alerts are consistent with user's expectation even with transmission delays. If any of the sources go offline or are slow to send sessions, ESA will pause to ensure that sessions with the same capture timestamps are correlated together.


For example, you have two sources with events that occur at 10:00 a.m. Using Capture Time Ordering, these events are held in the buffer until the ESA detects that all events occurring at 10:00 a.m. have been added to the buffer. Once all the events have arrived, events are then processed using EPL rules. This ensures that a rule has all events with the same time-stamp from different sources in order to obtain correct results. If, for example, one Concentrator lags behind another, the ESA pauses until it has all the events time-stamped at 10:00 a.m. from both sources before it runs the EPL rules against the events.

Caution: Although this feature increases accuracy, it impacts performance. The default configuration of the ESA ensures that data is constantly streaming, but because Capture Time Ordering uses a buffer, it takes longer to process events. This is especially true if the ESA must pause for any length of time to wait for the buffer to fill. There are several parameters you can configure (see below) to handle this situation; however, there may still be performance impact.

By default, this feature is disabled.

Capture Time Order Workflow

The following diagram shows the workflow when Capture Time Ordering is enabled.



1. Events are time-stamped as they are captured by the Decoder.
2. After Concentrator processing, events are buffered and ordered. The buffer size is calculated via two parameters **MaxEPSExpectedPerSource** (the maximum volume of traffic (EPS) you expect **per source** for the ESA to receive) times **TimeOrderHoldWindow** (the amount of time to allow for events to arrive from all sources).
3. The ordered events are then correctly correlated in EPL rules.

Prerequisites

Two or more Concentrators must be configured as a data source in ESA.


When the **StreamEnabled** parameter is set to true, it is important that all the machines running Core Services should be in NTP Sync.

Procedures

The following procedures tell you how to enable and configure Capture Time Ordering.

Enable Buffering and Capture Time Ordering

Note: After an upgrade or in a high EPS environment, you need to re-add datasources to start seeing the benefits. Or, you must wait until the sessions catch up before you enable Capture Time Ordering.

1. Go to **ADMIN > Services**, select your ESA service, and then select  > **View > Explore**.
2. Go to **Workflow > Source > nextgenAggregationSource**.
3. Set the **StreamEnabled** attribute to **true**. StreamEnabled allows ESA to buffer events received from Concentrators.
4. Set the **TimeOrdered** attribute to **true**. This enables the buffered events to be ordered by the time stamp from the Concentrator.

Configure Capture Time Ordering

When you work with Capture Time Ordering, you need to configure several other parameters to ensure performance. The following table shows parameters and their function. Configuring these parameters requires knowledge of your traffic volume and rate.


Note: If you do not know your traffic volume or latency, consult with your Professional Services representative before configuring this feature.

Parameter	Description
MaxEPSExpectedPerSource	<p>Specify the maximum volume of traffic (EPS, or events per second) you expect for the ESA service to receive from your busiest source (for example, if one source receives 20K EPS, and another receives 25K EPS, set the value at 25K EPS).</p> <p>If you set this rate too low, there is a short-term impact on performance. However, ESA automatically increases the value for MaxEPSExpectedPerSource as needed to make progress in Time Ordered mode.</p> <p>The default value is 20K.</p>
TimeOrderHoldWindow	<p>Specify in seconds (whole integers) the amount of time to allow for events to arrive from all sources.</p> <p>Configure this value based on the latency between the sources.</p> <p>The default value is 2 seconds. Decreasing this value can increase the chance of dropped events. Increasing this value can decrease performance because more memory is consumed.</p>
IdleSourceAdvanceAfterSeconds	<p>Specify the interval (in seconds) after which the ESA takes an idle source (no events are coming from the source, but the source is not offline) out of the equation to allow progress on a capture time ordered stream. The default value is 0, meaning that the ESA waits indefinitely for events to arrive.</p>
OfflineSourceAdvanceAfterSeconds	<p>Specify the interval (in seconds) after which the ESA takes an offline source out of the equation to allow progress on a capture time ordered stream. The default value is 0, which means the ESA waits indefinitely. This parameter does not affect the re-connection retries; those are performed in all cases.</p>

Troubleshooting Tips

Using this feature, it is possible to encounter a situation where events become backlogged. To fix this issue, you can perform one of the following options.


Disable Capture Time Ordering

1. Go to **ADMIN > Services**, select your ESA service, and then  > **View > Explore**.
2. Go to **Workflow > Source > nextgenAggregationSource**.
3. Set the `StreamEnabled` attribute to false.
4. Set the `TimeOrdered` attribute to false.

If you disable Capture Time Ordering, you will lose the backlogged data, and events will no longer be ordered by capture time.

Disable Position Tracking

Position tracking allows ESA to track where it stopped processing events if the ESA stops or is shut down. Position tracking is enabled by default with Capture Time Ordering. If you disable position tracking, this allows ESA to skip the backlogged events. For example, if the ESA goes down at 7:00 a.m., and you restart it at 11:00 a.m. with position tracking disabled, the ESA will start processing events that occurred at 10:55 a.m. With position tracking enabled, the ESA will start processing events at the point at which it stopped.

1. Go to **ADMIN > Services**, select your ESA service, and then select  > **View > Explore**.
2. Go to **Workflow > Source > nextgenAggregationSource**.
3. Set the `PositionTrackingEnabled` attribute to false.

If you disable Position Tracking, you will lose the backlogged data, but going forward, events will be ordered by capture time.


```
(\ 'CreateUser\ '))\n\tOR\n\t/* Statement: Instance state change */\n\t(medium = 32 AND device_type = \'awscloudtrail\' AND event_desc IN (\ 'TerminateInstances\' , \'RunInstances\ '))\n\t).win:time(310 seconds)\n\tMATCH_RECOGNIZE (\n\tPARTITION BY alias_host\n\tMEASURES E1 as e1_data , E2 as e2_data\n\tPATTERN (E1+ E2)\n\tDEFINE\n\tE1 as E1.medium = 32 AND E1.device_type = \'awscloudtrail\' AND E1.event_desc IN (\ 'CreateUser\ ')\n\tE2 as E2.medium = 32 AND E2.device_type = \'awscloudtrail\' AND E2.event_desc IN (\ 'TerminateInstances\' , \'RunInstances\ ')\n);\n\n" identity: "admin" userRole: "ROLE_ESA_ADMINISTRATOR"
```

- Update log example:** 2018-08-15 19:48:47,941 deviceVersion: "11.2.0.0-SNAPSHOT" deviceService: "EVENT_STREAM_ANALYSIS" category: SYSTEM operation: "**UPDATE RULE**" parameters: "Epl Module Identifier: esa000155, Esper Instance: default, Epl Rule Enabled: true, Trial Rule: true" key: "Epl Rule: /*\nVersion: 2\n*/\n\nmodule Module_esa000155;\n\n\n@Name(\ 'Module_esa000155_Alert\ ')\n@RSAAlert(oneInSeconds=0, identifiers={\"alias_host\"})\n\nSELECT * FROM Event(\n\t/* Statement: User permission change */\n\t(medium = 32 AND device_type = \'awscloudtrail\' AND event_desc IN (\ 'CreateUser\ '))\n\tOR\n\t/* Statement: Instance state change */\n\t(medium = 32 AND device_type = \'awscloudtrail\' AND event_desc IN (\ 'TerminateInstances\' , \'RunInstances\ '))\n\t).win:time(310 seconds)\n\tMATCH_RECOGNIZE (\n\tPARTITION BY alias_host\n\tMEASURES E1 as e1_data , E2 as e2_data\n\tPATTERN (E1+ E2)\n\tDEFINE\n\tE1 as E1.medium = 32 AND E1.device_type = \'awscloudtrail\' AND E1.event_desc IN (\ 'CreateUser\ ')\n\tE2 as E2.medium = 32 AND E2.device_type = \'awscloudtrail\' AND E2.event_desc IN (\ 'TerminateInstances\' , \'RunInstances\ ')\n);\n\n" identity: "admin" userRole: "ROLE_ESA_ADMINISTRATOR"
- Delete log example:** 2018-08-15 19:48:47,972 deviceVersion: "11.2.0.0-SNAPSHOT" deviceService: "EVENT_STREAM_ANALYSIS" category: SYSTEM operation: "**DELETE RULE**" parameters: "Epl Module Identifier: esa000155, Esper Instance: default, Epl Rule Enabled: true, Trial Rule: true" key: "Epl Rule: /*\nVersion: 2\n*/\n\nmodule Module_esa000155;\n\n\n@Name(\ 'Module_esa000155_Alert\ ')\n@RSAAlert(oneInSeconds=0, identifiers={\"alias_host\"})\n\nSELECT * FROM Event(\n\t/* Statement: User permission change */\n\t(medium = 32 AND device_type = \'awscloudtrail\' AND event_desc IN (\ 'CreateUser\ '))\n\tOR\n\t/* Statement: Instance state change */\n\t(medium = 32 AND device_type = \'awscloudtrail\' AND event_desc IN (\ 'TerminateInstances\' , \'RunInstances\ '))\n\t).win:time(310 seconds)\n\tMATCH_RECOGNIZE (\n\tPARTITION BY alias_host\n\tMEASURES E1 as e1_data , E2 as e2_data\n\tPATTERN (E1+ E2)\n\tDEFINE\n\tE1 as E1.medium = 32 AND E1.device_type = \'awscloudtrail\' AND E1.event_desc IN (\ 'CreateUser\ ')\n\tE2 as E2.medium = 32 AND E2.device_type = \'awscloudtrail\' AND E2.event_desc IN (\ 'TerminateInstances\' , \'RunInstances\ ')\n);\n\n" identity: "admin" userRole: "ROLE_ESA_ADMINISTRATOR"

Each log contains the following parameters:

- Time stamp: Time the rule was modified. Example: 2018-08-15 19:48:47,972
- DeviceVersion: Version of your ESA service. Example: "11.2.0.0-SNAPSHOT"
- DeviceService: Example: EVENT_STREAM_ANALYSIS
- Category: Example: SYSTEM
- Operation: Examples: CREATE RULE, UPDATE RULE, DELETE RULE
- Parameters: Placeholder for the following keys:
 - Epl Module Identifier: unique identifier for the rule. Example: esa000155
 - Esper Instance: Esper instance on which rule is deployed. Example: default
 - Epl Rule Enabled: Displays if the rule is enabled or not. Example: EPL Rule Enabled: false
 - Trial Rule: Displays if the rule is configured as a trial rule or not. Example: Trial Rule: true
 - Epl Rule: Displays the rule syntax. Example:

```
"Epl Rule: /*\nVersion: 2\n*/\n\nmodule Module_esa000155;\n\n\n@Name
(\'Module_esa000155_Alert\')\n@RSAAlert(oneInSeconds=0, identifiers=
{\\"alias_host\\"})\n\nSELECT * FROM Event(\n\t/* Statement: User
permission change */\n\t(medium = 32 AND device_type = \'awscloudtrail\'
AND event_desc IN (\'CreateUser\'))\n\tOR\n\t/* Statement: Instance state
change */\n\t(medium = 32 AND device_type = \'awscloudtrail\' AND event_
desc IN (\'TerminateInstances\' , \'RunInstances\'))\n\t).win:time(310
seconds)\n\tMATCH_RECOGNIZE (\n\tPARTITION BY alias_host\n\tMEASURES E1
as e1_data , E2 as e2_data\n\tPATTERN (E1+ E2)\n\tDEFINE\n\tE1 as
E1.medium = 32 AND E1.device_type = \'awscloudtrail\' AND E1.event_desc
IN (\'CreateUser\'),\n\tE2 as E2.medium = 32 AND E2.device_type =
\'awscloudtrail\' AND E2.event_desc IN (\'TerminateInstances\' ,
\'RunInstances\'))\n); \n\n"
```

- Identity: Example: "admin"
- userRole: Example: "ROLE_ESA_ADMINISTRATOR"

Note: When a rule is disabled, two logs are generated for the same rule. First a 'Delete Rule' [Rule enabled attribute = true] audit log is created, followed by a 'Create Rule' [Rule enabled attribute =false] audit log.

Verify ESA Server Version

1. Use ssh to connect to the ESA service and log in as the root user.
2. Type the following command and press **ENTER**:

```
rpm -qa | grep rsa-nw-esa-server
```

 The ESA server version is displayed.

References

This section is a collection of references, which describe the user interface for ESA Configuration in NetWitness Platform.

See the following topics for details:

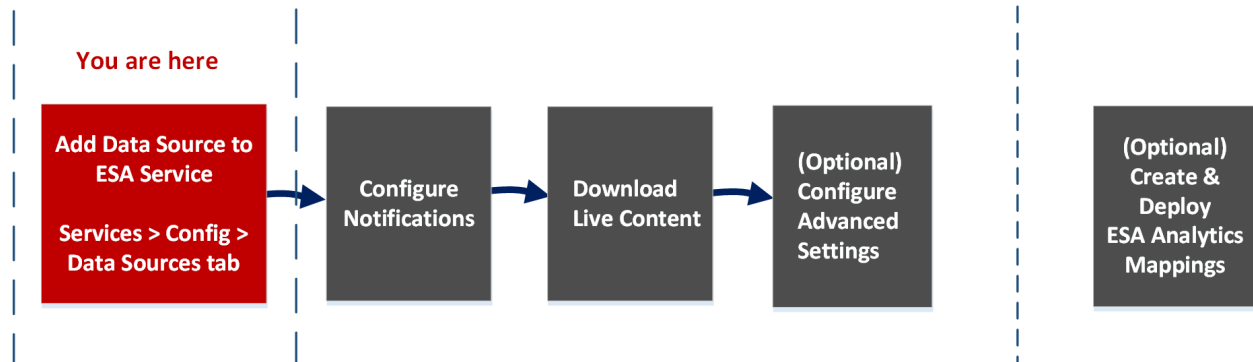
- [Services Config View Advanced Tab](#)
- [Services Config View Data Sources Tab](#)
- [ESA Analytics Mappings](#)
- [Module Settings](#)
- [Whois Lookup Service Configuration](#)

Services Config View Data Sources Tab

The **Services Config view > Data Sources** tab of an ESA service enables you to configure the sources that ESA uses to analyze data. An ESA service ingests data from Concentrators to detect incidents and alert analysts to potential threats.

Workflow

This workflow shows the overall process for configuring ESA. It also shows where configuring data sources is located in the process.



ESA has two services, the Event Stream Analysis service (ESA Correlation Rules) and the Event Stream Analytics Server service (ESA Analytics). The first four procedures shown pertain to configuring the Event Stream Analysis service:

- **Add Data Source to ESA Service ***
- Configure Notifications
- Download Live Content
- (Optional) Configure Advanced Settings

The last procedure is separate from the others and pertains to creating mappings for the ESA Analytics services to start automatically detecting advanced threats:

- (Optional) Create and Deploy ESA Analytics Mappings

What do you want to do?

Role	I want to ...	Show me how
Administrator	Add a Concentrator as a data source to the Event Stream Analysis Service *	See Configure ESA Correlation Rules and Step 1. Add a Data Source to an ESA Service
Administrator	Configure Notifications	See "Notification Methods" in the <i>Alerting with ESA Correlation Rules User Guide</i> .
Administrator	Download Live Content	See "Download Configurable RSA Live Rules" in the <i>Alerting with ESA Correlation Rules User Guide</i> .
Administrator	Configure Advanced Settings	See Step 2. Configure Advanced Settings for an ESA Service

*You can complete these tasks here (that is in the Services Config view Data Sources tab).


Related Topics

- See "Add or Update a Host" in the *Host and Services Getting Started Guide*

Quick Look

To access the Data Sources tab, go to **ADMIN > Services >** (Select an ESA service) >   > **View > Config.**

The following figure shows the Services Config view Data Sources tab for an ESA service.



Toolbar

The following table describes the options in the toolbar.

Option	Description
	Adds a new data source to the ESA service.
	Deletes a data source from the ESA service.
	Edits a data source. You must have the username and password credentials for the service in order to make changes.
Enable	Enables the selected data source.
Disable	Disables the selected data source.

Data Sources

The Data Sources list shows all of the data sources added to the ESA service. The following table describes the columns the Data Sources list.

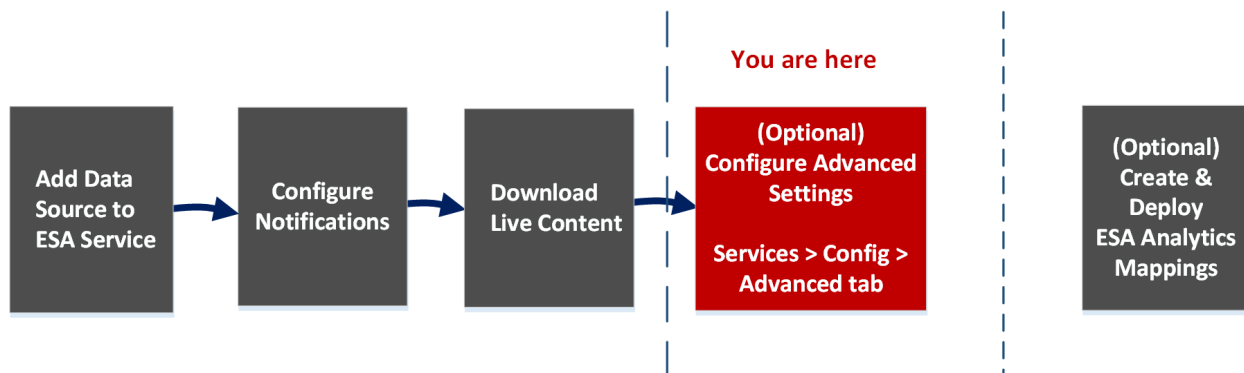
Column	Description
Name	The name of the data source service.
Address	The address of the data source service.
Port	The port used by the data source.
User	The user connected with the data source.
Enabled	Indicates if the data source is enabled.
SSL	Indicates if SSL communication is enabled.
Compression	Indicates if compression is enabled.

Services Config View Advanced Tab

The **Services Config view > Advanced** tab of an ESA service enables you to configure advanced settings. In the Advanced view, you can configure advanced settings to improve performance, to preserve events for rules with multiple events, to buffer events in memory, and to set the number of events to be stored on the ESA.

Workflow

This workflow shows the overall process for configuring ESA. It also shows where configuring advanced settings is located in the process.



ESA has two services, the Event Stream Analysis service (ESA Correlation Rules) and the Event Stream Analytics Server service (ESA Analytics). The first four procedures shown pertain to configuring the Event Stream Analysis service:

- Add Data Source to ESA Service
- Configure Notifications
- Download Live Content
- **(Optional) Configure Advanced Settings ***

The last procedure is separate from the others and pertains to creating mappings for the ESA Analytics services to start automatically detecting advanced threats:

- (Optional) Create and Deploy ESA Analytics Mappings

What do you want to do?

Role	I want to ...	Show me how
Administrator	Add a Concentrator as a data source to the Event Stream Analysis Service	See Configure ESA Correlation Rules and Step 1. Add a Data Source to an ESA Service .
Administrator	Configure Notifications	See "Notification Methods" in the <i>Alerting with ESA Correlation Rules User Guide</i> .
Administrator	Download Live Content	See "Download Configurable RSA Live Rules" in the <i>Alerting with ESA Correlation Rules User Guide</i> .
Administrator	Configure Advanced Settings *	See Step 2. Configure Advanced Settings for an ESA Service .

*You can complete these tasks here (that is in the Services Config view Advanced tab).


Related Topics

- See "Add or Update a Host" in the *Host and Services Getting Started Guide*

Quick Look

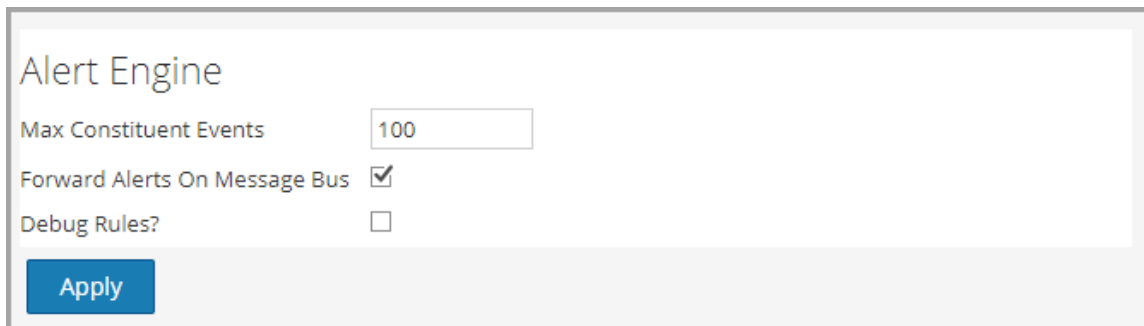
To access the Advanced tab, go to **ADMIN > Services >** (Select an ESA service) >  > **View > Config**.

The following figure shows the Services Config view Advanced tab for an ESA service.



Alert Engine Settings

In the Alert Engine section, you specify values to preserve events for rules that choose multiple events. The following figure shows the Alert Engine section.



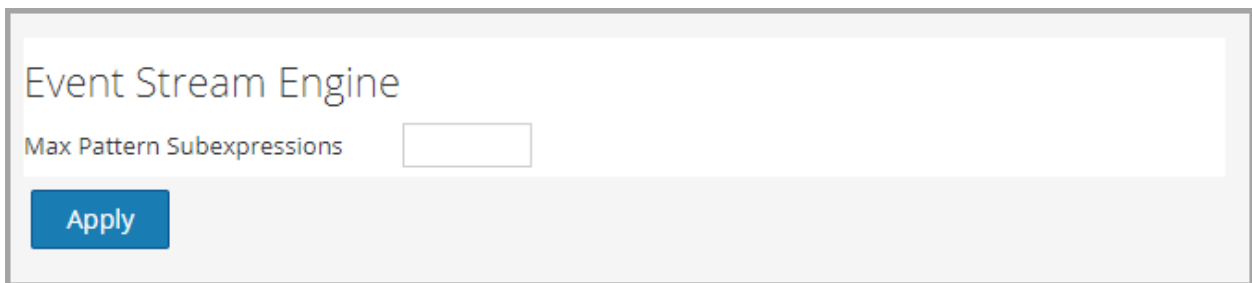
The following table lists the parameters in the Alert Engine section and their descriptions.

Parameter	Description
Max Constituent Events	For rules that contain multiple events, this configuration value determines how many of the associated events are preserved. For example, if a rule fires an alert with 200 associated events and this parameter is set to 100, only the first 100 are preserved by ESA, the rest are dropped. The default value is 100.

Parameter	Description
Forward Alerts On Message Bus	To forward ESA alerts for NetWitness Respond, you must select this option. The ESA alerts generated will be sent to the Message Bus and subsequently to Respond. This option is selected by default. You may want to ensure that the Respond Server service is running.
Debug Rules?	Selecting enables debugging rules.

Event Stream Engine Settings

In the Event Stream Engine section, you specify details to improve performance. The following figure shows the Event Stream Engine section.



The following table lists the parameter in the Event Stream Engine section and its description.

Parameter	Description
Max Pattern Subexpressions	Certain rules require ESPER to maintain subexpressions in memory before deciding to fire them or not. These subexpressions consume memory and if left unchecked may cause the service to go down with memory exhaustion. This parameter is a safety measure that keeps such memory hogging rules under check. If a rule exceeds the specified number of subexpressions, its processing is delayed. The default value is 0, which means this setting is disabled. You must set a value if there are service stability issues.

Whois Lookup Service Configuration

In the Whois Lookup Configuration panel (ADMIN > System > Whois), you configure a connection to the Whois Lookup service for your preconfigured ESA Analytics modules used in RSA Automated Threat Detection. The Whois Service enables you to get accurate data about domains that you connect to. In order to ensure effective scoring, it is important that you configure the Whois service settings.

You must have an RSA Live account to use this service.

If you configured a Live account in the Live Services panel (ADMIN > System > Live Services), the Whois Lookup Service is automatically configured for you. You just need to check the connection of the Whois Lookup service.

Note: If you do not have an RSA Live account, you can create one at the RSA Live Registration Portal:
<https://cms.netwitness.com/registration/>
 The *Live Services Management Guide* provides additional information.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Configure the Whois Lookup service.	Configure the Whois Lookup Service
Administrator	Check the connection of the Whois Lookup service.	Configure the Whois Lookup Service

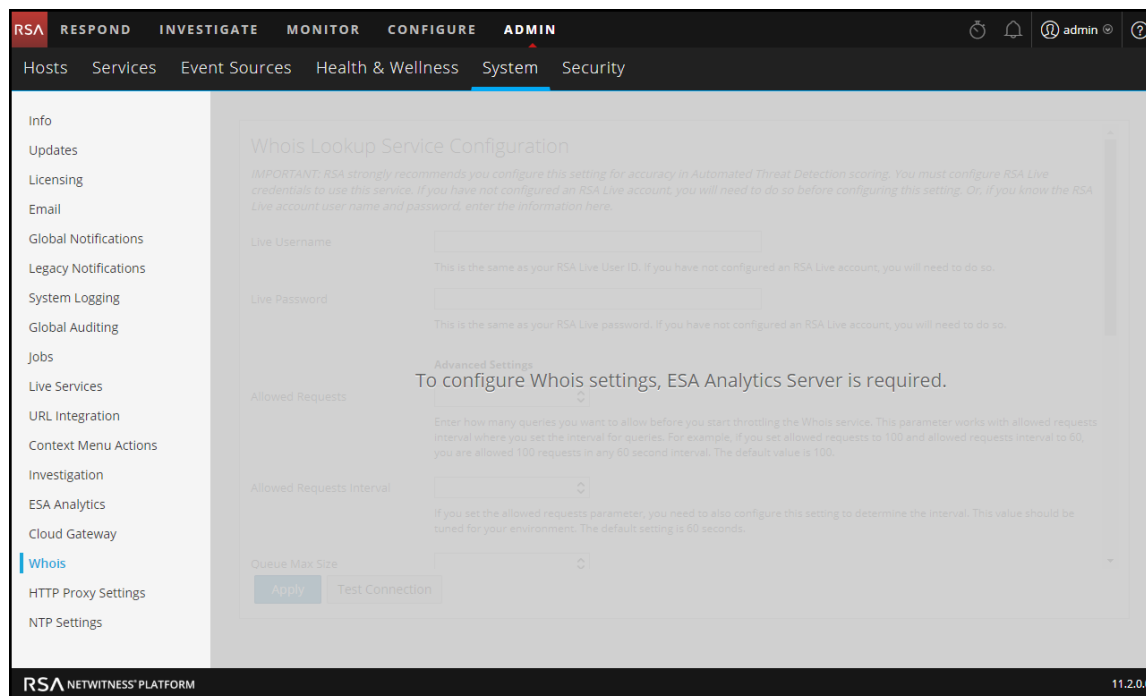
Related Topics

- [ESA Analytics Mappings](#)

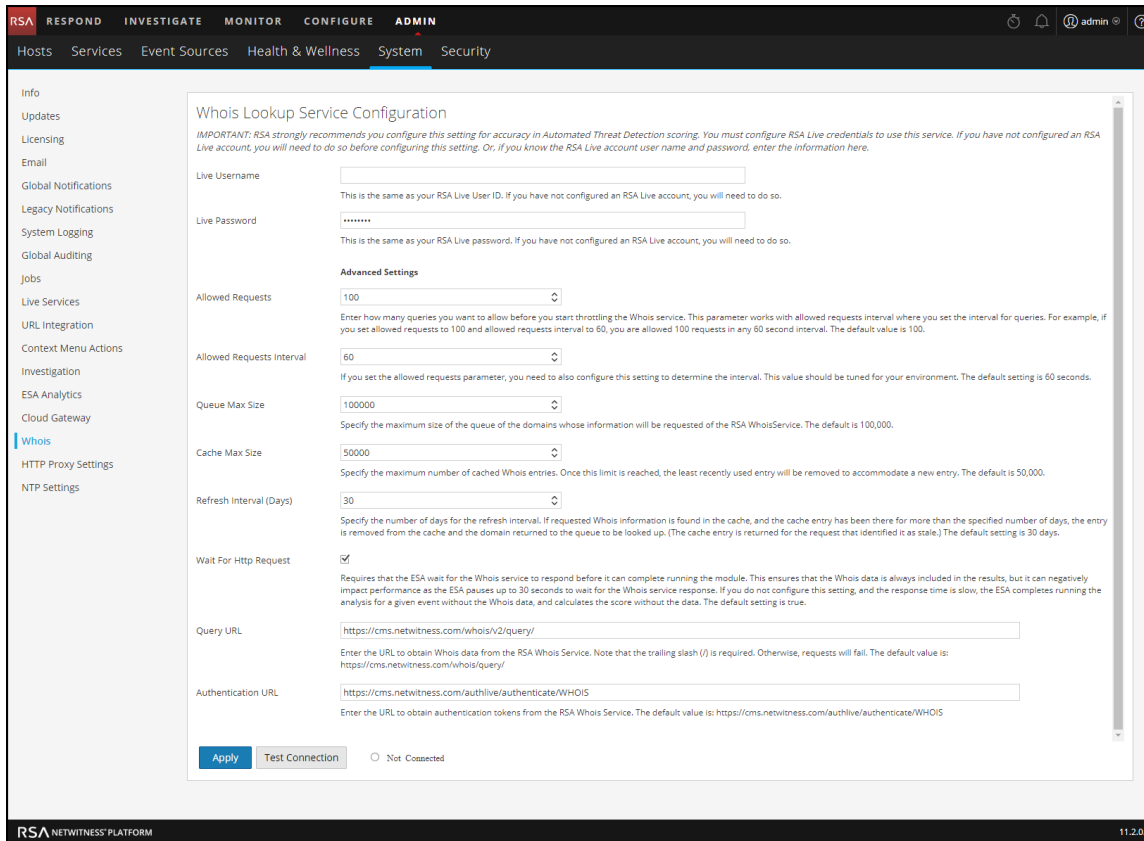
Quick Look

To access the Whois Lookup Service Configuration, go to ADMIN > System and in the options panel, select Whois.

The ESA Analytics Server service must be available (shows a green circle) in the ADMIN > Services view. If you do not have an ESA Analytics Server service available, you will see the following panel.



If you have an ESA Analytics Server service available, you will see the following panel.



The following table describes the listed Whois Lookup Service configuration settings.

Parameter	Description
Live Username	Required only if you did not already configure the Whois Lookup service. Enter the authentication credential for the RSA Whois Server. This is the same as your RSA Live User ID. If you have not configured an RSA Live account, you will need to do so. The default value is "whois."
Live Password	Required only if you did already configure the Whois Lookup service. Enter the authentication credential for the RSA Whois Server. This is the same as your RSA Live password. If you have not configured an RSA Live account, you will need to do so. The default value is null.
Allowed Requests	(Optional) Enter how many queries you want to allow before you start throttling the Whois service. This parameter works with Allowed Requests Interval (in seconds), where you set the interval for queries. For example, if you set Allowed Requests to 100 and Allowed Requests Interval to 60, you are allowed 100 requests in any 60 second interval. The default value is 100.

Parameter	Description
Allowed Requests Interval	<p>(Optional) If you set the Allowed Requests parameter, you need to also configure this setting to determine the interval. This value should be tuned for your environment.</p> <p>The default setting is 60 seconds.</p>
Queue Max Size	<p>(Optional) Specify the maximum size of the queue of the domains whose information will be requested of the RSA WhoisService.</p> <p>The default is 100,000.</p>
Cache Max Size	<p>(Optional) Specify the maximum number of cached Whois entries. Once this limit is reached, the least recently used entry will be removed to accommodate a new entry.</p> <p>The default is 50,000.</p>
Refresh Interval (Days)	<p>(Optional) Specify the number of days for the refresh interval. If requested Whois information is found in the cache, and the cache entry has been there for more than the specified number of days, the entry is removed from the cache and the domain returned to the queue to be looked up. (The cache entry is returned for the request that identified it as stale.)</p> <p>The default setting is 30 days.</p>
Wait For HTTP Request	<p>(Optional) Requires that the ESA wait for the Whois service to respond before it can complete running the module. This ensures that the Whois data is always included in the results, but it can negatively impact performance as the ESA pauses up to 30 seconds to wait for the Whois service response.</p> <p>If you do not configure this setting, and the response time is slow, the ESA completes running the analysis for a given event without the Whois data, and calculates the score without the data.</p> <p>The default setting is true.</p>
Query URL	<p>(Optional) Enter the URL to obtain Whois data from the RSA Whois service. The trailing slash (/) is required. Otherwise, requests will fail.</p> <p>The default value is: <code>https://cms.netwitness.com/whois/v2/query/</code></p>
Authentication URL	<p>(Optional) Enter the URL to obtain authentication tokens from the RSA Whois service. The default value is: <code>https://cms.netwitness.com/authlive/authenticate/WHOIS</code></p>

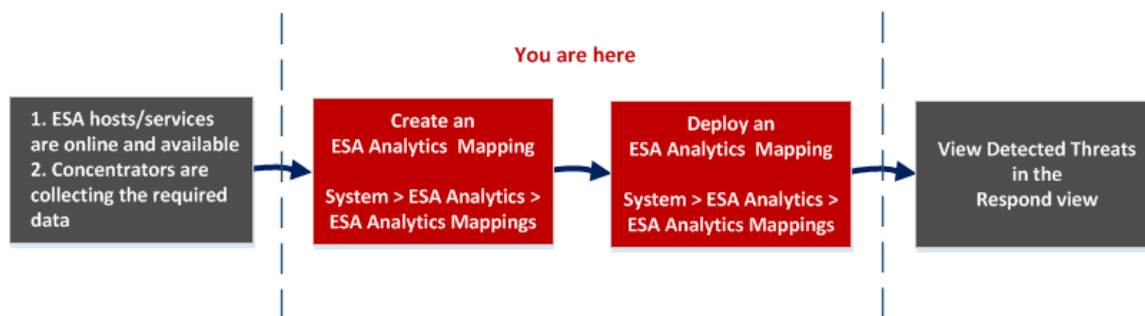
ESA Analytics Mappings

In the ESA Analytics Mappings panel (ADMIN > System > ESA Analytics), you define how the RSA Automated Threat Detection functionality should automatically detect advanced threats. You can analyze the data that resides on one or more Concentrators by selecting a preconfigured ESA Analytics module.

To better utilize your network resources and reduce unnecessary data flow, you can map multiple data sources, such as Concentrators, to available ESA Analytics services in order to process data more efficiently and take advantage of additional capacity.

Workflow

This workflow shows the process for creating and enabling an ESA Analytics mapping to start automatically detecting advanced threats.



Before you create an ESA Analytics mapping, ensure that the ESA hosts and services that you want to use for your mappings are online and available. All of the services need to be in sync with a consistent time source. Also ensure that the Concentrators are collecting the required data. When you create an ESA Analytics mapping, you select an ESA Analytics module to map, such as Suspicious Domains. Then you select the data sources, such as Concentrators, to use for that module along with an ESA Analytics service to process the data. When you are ready to start aggregating data, you deploy the mapping. Analysts can view detected threats for that module in the Respond view.

What do you want to do?

Role	I want to ...	Show me how
Administrator	Verify that the ESA hosts and services are online and available.	ADMIN > HOSTS and ADMIN > SERVICES See <i>Hosts and Services Getting Started Guide</i> .
Administrator	Ensure that the Concentrators are collecting the required data.	See <i>Broker and Concentrator Configuration Guide</i>
Administrator	Create ESA Analytics mappings.*	Mapping ESA Data Sources to Analytics Modules
Administrator	Deploy ESA Analytics mappings.*	Mapping ESA Data Sources to Analytics Modules
Administrator, Analyst	View detected threats.	See <i>NetWitness Respond User Guide</i> .


*You can complete these tasks here (that is in the ESA Analytics Mappings panel).

Related Topics

- [Configure ESA Analytics](#)
- [Update a Mapping](#)
- [Undeploy a Mapping](#)
- [Delete a Mapping](#)
- [Change the Warm-up Period and Lag Time](#)
- [Module Settings](#)

Quick Look



The following example illustrates an ESA Analytics mapping. The configuration defines the data sources for the selected module and the ESA Analytics service that will process the events from those data sources.



- 1 Displays the ESA Analytics Mappings panel.
- 2 Shows the status of the ESA Analytics mapping.
- 3 The name of the module that is mapped.
- 4 Data sources, such as Concentrators, assigned to the mapping.
- 5 ESA Analytics service that processes the data for the mapping.
- 6 Warm-up period configuration (in hours) on the data sources for the mapping.
- 7 Lag configuration (in minutes) on the data sources for the mapping.
- 8 Actions for changing module settings, deploying module mappings, and undeploying module mappings.

Toolbar


The following table describes the toolbar actions.

Icon / Button	Description
	Opens the Create Mappings dialog where you can create an ESA Analytics mapping. Create a separate mapping for each module. After creating and reviewing the mappings, you deploy them.
	Deletes an ESA Analytics Mapping. <ul style="list-style-type: none"> You can delete a mapping with a status of Undeployed at any time. Since a mapping in the Undeployed state is not deployed and is not running, it does not affect data aggregation. Deleting a deployed mapping clears the configuration on the ESA server, reverts the deployment for that mapping, and stops pulling data from the data source for that module. You should undeploy a mapping with a status of Deployed before deleting it.
Deploy Now	After you create your mappings, you need to deploy them in order to start aggregating data for the modules. You can select one or more mappings with a status of Undeployed to deploy.


Note: If you want to make changes to a deployed mapping, such as adding or removing Concentrators or changing the service, you must undeploy and delete the existing mapping and then create and deploy a new mapping for that module.

ESA Analytics Mappings

The following table describes the listed ESA Analytics mappings.

Title	Description
	To select an individual mapping, select the checkbox next to the mapping.
Status	Shows the status of the mapping. There are two statuses: <p>Undeployed - An undeployed mapping maps an ESA Analytics module to sources and an ESA Analytics service. It does not start aggregating data for the module until you deploy the mapping.</p> <p>Deployed - A deployed mapping is deployed and running. In a deployed mapping, the selected ESA Analytics service uses query-based aggregation to collect the appropriate filtered events for the selected module from the Concentrators.</p>
Module	Indicates the selected ESA Analytics module. An ESA Analytics module is a pipeline composed of activity objects that enrich an event with additional information through mathematical computations. The module resides within the ESA Analytics service.

Title	Description
Sources	Sources are the data sources, such as Concentrators, from which ESA will aggregate the data for the specified module.
Service	Indicates the ESA Analytics service that will process the data for the specified module. The selected service needs to be in sync with a consistent time source.
Warm-Up Period (Hours)	<p>Specifies a warm-up duration (in hours). A warm-up period is required to allow Automated Threat Detection to "learn" your traffic. The warm-up period should run when typical traffic is running. During this time, alerting for your module mapping is suppressed. The Warm-up Period primes the module with historical data and guarantees that the specified number of hours of data collection completes before sending alerts.</p> <p>RSA provides preconfigured ESA Analytics modules. Each module type has a default warm-up period defined, which you can adjust to your environment, if necessary. After this warm-up period, alerts can be viewed.</p> <p>For more information about Warm-up Period and Lag time, see Module Settings.</p>
Lag Time (Minutes)	<p>Specifies a constant time delay in minutes, which is added to avoid losing events being processed by the data sources during periods of heavy activity. For example, Concentrator performance varies depending on factors such as incoming load, ongoing queries, and indexing. Due to these factors, a Concentrator may not aggregate events in real-time, which leads to the delay.</p> <p>The Lag parameter gives the Concentrator a chance to finish aggregating all of the data. After the warm-up period completes, data aggregation continues at Current (System) Time - Lag Time. This is useful when a Concentrator is slow in aggregating data. The Lag time guarantees that the module does not process data that arrives to the Concentrator within the Lag time window so there is adequate delay to ensure all events that get generated in the enterprise can be processed by the module.</p> <p>For example, if Lag time is 30 minutes, and the current time is 2:00 PM, the Concentrator starts pulling records at 1:30 PM. The Lag time window, 30 minutes in this example, remains constant as time advances. When the current time advances to 2:01 PM, the Concentrator pulls the next minute of data at 1:31 PM, and so on.</p> <p>Important: The Lag time defines the buffer between the current time and the time when the module ingests the data.</p> <div style="border: 1px solid yellow; padding: 5px; margin-top: 10px;"> <p>Caution: RSA recommends that Administrators adjust the Lag parameter dynamically based on the performance of each of the individual Concentrators to avoid missing any events during aggregation.</p> </div> <p>For more information about Warm-up Period and Lag time, see Module Settings.</p>

Title	Description
	<p>Enables you to select additional actions for the selected module mapping:</p> <ul style="list-style-type: none"> • Edit Module - Enables you to configure the warm-up period and lag time for the selected module mapping. • Deploy - Deploys the selected module mapping. The specified ESA Analytics service starts pulling data from the data sources for that module. • Undeploy - Undeploys the selected module mapping. The specified ESA Analytics service stops pulling data from the data sources for that module. <div style="border: 1px solid yellow; padding: 5px; margin-top: 10px;"> <p>Caution: Undeploying a mapping with a status of Deployed will affect data aggregation for that module.</p> </div>

Module Settings

After you create or deploy a module mapping in the ESA Analytics Mappings panel (ADMIN > System > ESA Analytics), you have the option to change some module configurations for that mapping.

What do you want to do?


Role	I want to ...	Show me how
Administrator	Change the warm-up period for an undeployed module mapping.	Change the Warm-up Period and Lag Time
Administrator	Change the warm-up period for a module mapping during the warm-up period.	Change the Warm-up Period and Lag Time
Administrator	Change the warm-up period for a module mapping after the warm-up period is complete.	Change the Warm-up Period and Lag Time

Related Topics

- [Mapping ESA Data Sources to Analytics Modules](#)
- [ESA Analytics Mappings](#)

Quick Look

To access the module settings, in the ESA Analytics Mappings panel, select the mapping that you want to change and in the **Actions** column, select  > **Edit Module**. The Module Settings dialog has a Configurations section and a Warm-Up State section.



Configurations

The Configurations section enables you to change the Warm-Up Period and Lag Time configurations. The following table describes the settings available for an ESA Analytics module mapping.

Field	Description
Module	Shows the name of the mapped module.
Service	Shows the ESA Analytics service that processes the data for the mapping.
Sources	Shows the mapped data sources and the URLs used to communicate with ESA.
Warm-Up Period (Hours)	<p>Specifies a warm-up duration in hours. A warm-up period is required to allow Automated Threat Detection to "learn" your traffic. The warm-up period should run when typical traffic is running. During this time, alerting for your module mapping is suppressed. The Warm-up Period primes the module with historical data and guarantees that the specified number of hours of data collection completes before sending alerts.</p> <p>RSA provides preconfigured ESA Analytics modules. Each module type has a default warm-up period defined, which you can adjust to your environment, if necessary. After this warm-up period, alerts can be viewed.</p> <p>You can update the Warm-Up Period of a deployed module mapping depending on whether or not the warm-up period is complete:</p> <ul style="list-style-type: none">• During the warm up period - You can add hours to the warm-up period or subtract any remaining warm-up time.• The warm-up period is complete - You can add hours to the warm-up period by adding the difference between the current time and the First Event Time to the hours that you want to add. <p>For example, a warm-up period of 10 hours is complete and the First Event Time shows 12:00:00. The current (system) time is 16:00:00 (4 hours later) and you want to add 5 more hours to the warm-up time. To do this, you need to add 9 hours (4+5=9) to the warm-up period of 10, so you would set the new warm-up period to 19 hours.</p> <p>You cannot decrease the warm-up period if it is complete, unless you delete the mapping and create a new one.</p> <p>The Warm-up Period value is specific to a particular mapping and it applies to all Concentrators within that mapping after you deploy it. If a Concentrator is shared between two modules with different warm-up times, the Concentrator uses separate Warm-up Period values for each module mapping.</p>

Field	Description
Lag Time (Minutes)	<p>Specifies a constant time delay in minutes, which is added to avoid losing events being processed by the data sources during periods of heavy activity. For example, Concentrator performance varies depending on factors such as incoming load, ongoing queries, and indexing. Due to these factors, a Concentrator may not aggregate events in real-time, which leads to the delay.</p> <p>The Lag parameter gives the Concentrator a chance to finish aggregating all of the data. When you specify a Lag time, the first time the module deploys, data aggregation starts at Current (System) Time - Lag Time - Warm-Up Time. For example, if the current time is 2:00 PM, Lag time is 30 minutes, and Warm-up time is 4 hours, when the module deploys for the first time, data collection starts at 9:30 AM (2:00 PM - .5 hour - 4 hours).</p> <p>After the warm-up period completes, data aggregation continues at Current (System) Time - Lag Time. This is useful when a Concentrator is slow in aggregating data. The Lag time guarantees that the module does not process data that arrives to the Concentrator within the Lag time window so there is adequate delay to ensure all events that get generated in the enterprise can be processed by the module.</p> <p>For example, if Lag time is 30 minutes, and the current time is 2:00 PM, the Concentrator starts pulling records at 1:30 PM. The Lag time window, 30 minutes in this example, remains constant as time advances. When the current time advances to 2:01 PM, the Concentrator pulls the next minute of data at 1:31 PM, and so on.</p> <p>Important: The Lag time defines the buffer between the current time and the time when the module ingests the data.</p> <p>The Lag time value is specific to a particular mapping and it applies to all Concentrators within that mapping after you deploy it. If a Concentrator is shared between two modules with different Lag times, the Concentrator uses separate Lag values for each module mapping.</p> <div style="border: 1px solid yellow; padding: 5px;"> <p>Caution: RSA recommends that Administrators adjust the Lag parameter dynamically based on the performance of each of the individual Concentrators to avoid missing any events during aggregation.</p> </div> <p>To determine the correct Lag Time, add together the following to get an environmental lag time:</p> <ol style="list-style-type: none"> 1. Log or Packet Latency - This is the time it takes for the Log Decoder to receive the logs or the (Packet) Decoder to receive packets. For example, the Log Decoder may get logs every 20 minutes. In this case, you would want to set Lag time to at least 20 minutes, preferably 25 minutes, so that you do not miss events. 2. Aggregation Latency - This is the time it takes to get the data from the Log Decoder to the Concentrator. 3. Other Buffer - Add in any additional time delay specific to your environment.

Warm-Up State

The Warm-Up State section provides information about the warm-up state, which you can use to determine the appropriate adjustments to the warm-up period.

Field	Description
Warmup Started At	The time when the first event was processed by the ESA Analytics module from the data source.
First Event Time	The time that the first event occurred. The warm-up time is based on this time.
Latest Event Time	The time that the latest event occurred.
Remaining Warm-Up Time	The number of hours remaining in the warm-up period.
Is Completed?	Indicates whether the warm-up period is complete. If it is true, the warm-up period is complete. If it is false, the module is still warming up and you can view the number of hours remaining in the Remaining Warm Up Time field.

