Micro Focus Security ArcSight Fusion in the ArcSight Platform User's Guide

Software Version: 1.6.0

Fusion in the ArcSight Platform User's Guide

Legal Notices

Micro Focus
The Lawn
22-30 Old Bath Road
Newbury, Berkshire RG14 1QN
UK

https://www.microfocus.com

Copyright Notice

© Copyright 2023 Micro Focus or one of its affiliates

Confidential computer software. Valid license from Micro Focus required for possession, use or copying. The information contained herein is subject to change without notice.

The only warranties for Micro Focus products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein.

No portion of this product's documentation may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's internal use, without the express written permission of Micro Focus.

Notwithstanding anything to the contrary in your license agreement for Micro Focus ArcSight software, you may reverse engineer and modify certain open source components of the software in accordance with the license terms for those particular components. See below for the applicable terms.

U.S. Governmental Rights. For purposes of your license to Micro Focus ArcSight software, "commercial computer software" is defined at FAR 2.101. If acquired by or on behalf of a civilian agency, the U.S. Government acquires this commercial computer software and/or commercial computer software documentation and other technical data subject to the terms of the Agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the Federal Acquisition Regulation ("FAR") and its successors. If acquired by or on behalf of any agency within the Department of Defense ("DOD"), the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of the Agreement as specified in 48 C.F.R. 227.7202-3 of the DOD FAR Supplement ("DFARS") and its successors. This U.S. Government Rights Section 18.11 is in lieu of, and supersedes, any other FAR, DFARS, or other clause or provision that addresses government rights in computer software or technical data.

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S.Government rights, patent policy, and FIPS compliance, see https://www.microfocus.com/about/legal/.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document home page of this Help contains the following identifying information:

- Software Version number
- Document Release Date, which changes each time the document is updated
- Software Release Date, which indicates the release date of this version of the software

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

https://www.microfocus.com/support-and-services/documentation

Welcome to Fusion

This User's Guide provides concepts, use cases, and contextual help for the **Fusion** common layer of services, helping you with the following activity:

- Manage users and groups of users
- Access the Reports Portal to create visuals and reports for analyzing data
- Search for alerts and events
- Check the integrity of your data
- Hunt for threats and vulnerabilities with built-in reports
- Respond to Threats
- Manage your MSSP profile and submit monthly EPS usage reports
- Manage ingested and imported data
- Create and use ArcSight dashboards to analyze data
- Access ArcMC to manage and monitor components in the ArcSight infrastructure
- Manage your user profile

About this Guide

This User's Guide provides concepts, use cases, and contextual help for the Fusion common layer of services.

Intended Audience

This book provides information for individuals who need to create users, groups, and roles; create and run reports and dashboards, and use the ArcSight Dashboard. These individuals have experience using security and identity management products, as well as creating reports and dashboards.

Additional Documentation

This documentation library includes the following resources

- *Quick Start Guide for Administrators*, which provides an overview of the products deployed in this suite and their latest features or updates
- User Guides and Release Notes for the capabilities that deployed in your ArcSight SaaS environment

For the most recent version of this guide and other ArcSight documentation resources, visit the documentation site for ArcSight SIEM as a Service.

Contact Information

We want to hear your comments and suggestions about this book and the other documentation included with this product. You can use the comment on this topic link at the bottom of each page of the online documentation, or send an email to Documentation-Feedback@microfocus.com.

For specific product issues, contact Micro Focus Customer Care.

Managing Users

You can add users or groups of users; create roles; and assign permissions to the roles for users and groups. There are default roles with appropriate permissions to use the product. If you manage groups, you can view their assigned permissions, roles, and users.

If you have the *Manage Roles* permission, you can change the permissions of any role assigned to your account except for those of the *System Admin*.

Managing Users and Groups of Users

You must have the appropriate permissions *to perform these functions.*

Click **ADMIN** > **Users and Groups**.

To delegate responsibility of managing large numbers of users across multiple managers, you can create groups. You can assign one or more managers to a group of users. Then managers assign roles to users in their groups.

Import Users from ArcSight Enterprise Security Manager

This function is not available in a SaaS environment.

To help you get started, you can import users already authorized for ESM. You need to have at least one role available in Fusion to assign to these users.

Click **ADMIN** > **Users and Groups** > group_name.

For more information, see the *Administrator Guide for ArcSight Platform*.

View Details of a Group

You can view the details of a group. As the manager of a group, you can also modify the group's settings.

1. Click **ADMIN** > **Users and Groups** > group_name.

2. (Conditional) As a manager of the group, you can also perform the following actions:

   - Add or remove users from this or other groups

   - Assign or remove roles for users in the current group

   - Add or remove managers from the current group

     > If you have the *System Admin* role, you can add and remove managers regardless whether you manage the group.

Create a New Group

1. Click **ADMIN** > **Users and Groups** > **Create Group**.

2. Specify a name for the group, then press **Enter**.

3. To manage the new group, perform the following actions:

   - Add users to this group

   - Assign roles to the users in this group

   - Add managers to this group

Create a New User

Users must have at least one role to ensure that they can log in.

1. Click **ADMIN** > **Users and Groups** > **Create User**.

2. Specify the email ID and name of the user.

3. Select the groups to which you want to add the user.

4. Select the roles that you want to grant to the user.

5. Click **Save**.

6. (Conditional) In a non-SaaS environment, specify the user's password.

> If SMTP is configured, the system notifies the new user over email to set up a password.

View a User's Profile

The user profile provides basic details about the user. If you are a manager of the user's account group, you can modify the user's account. You must also have appropriate permissions to make the modifications.

1. To find the user, perform one of the following actions:

   - Click **ADMIN** > **Users and Groups** > **Search Users**.

   - Click **ADMIN** > **Users and Groups** > group_name.

2. Select the user that you want to view.

3. (Optional) Modify the user's profile in one of the following ways:

   - Reset the password

   - Activate or deactivate the user

   - Change roles or permissions

   - Change group assignments

Change the User's Password

This function is not available in a SaaS environment.

*You must have the **Change User Password** permission, and be a manager of the user's account group.*

When you reset a user's password, the user receives a notification email automatically. The email does not include the new password. You must provide the new password to the user directly.

> In a SaaS environment, administrators and managers cannot create or change a user's password. Users can specify and reset their passwords by using the Advanced Authentication service.

1. Select **Users and Groups** > **Search Users**.
2. Select the user that you just created.
3. Click **RESET PASSWORD**.
4. Enter the password.
5. Click **SAVE**.
6. Notify the user of the new password.

Change the User's Status

*You must have the **Activate/Deactivate Users** permission, and be a manager of the user's account group.*

While you cannot delete a user, you can deactivate their account to prevent them from logging in to the system.

1. Adjust the **User Status** toggle switch to indicate Active or Inactive, as needed.
2. Click **SAVE**.

Change the User's Roles

*You must have the **Assign Roles to Users** permission, and be a manager of the user's account group.*

You can only assign those roles that you currently have. However, if you have the *Manage Groups* permission, you can assign any role to these users.

1. In the user's profile, select **Roles & Permissions**.

2. Select **Assign/Remove Roles**.

3. Change the user's roles, then select **Save**.

   Each role has a defined set of permissions. To change a user's permissions, you must change the assigned role or the permissions associated with a role.

Change the User's Group Assignments

*You must have the **Assign Users to Groups** permission, and a manager of the user's account group.*

Unless you have the *Manage Groups* permission, you can only assign those groups in which you are currently a member.

1. In the user's profile, select **Groups**.
2. Select **Add/Remove**.
3. Change the user's group assignments.

Assigning Permissions to Roles

You must have the appropriate permissions to create roles and assign permissions.

Click **ADMIN** > **Roles and Permissions**.

Fusion provides a set of roles that you can assign to your users. You can also create new roles with any combination of the available permissions. You can assign only the permissions and roles that you have yourself.

Available Permissions

Some permissions are available for any deployed product. Other permissions depend on the capabilities that you have deployed.

- Reports Permissions
- User Management Permissions
- ArcSight Permissions

Reports Permissions

The following table lists the permissions available when you add the Reports feature.

Function	Permissions	In the Reports Portal, allows users to...
Reports	Report Admin	View dashboards and reports Create subfolders Schedule reports Create data worksheets, dashboards, and reports View Admin reports Manage the data source (*not available in a SaaS environment*)
Reports	Design Reports	View dashboards and reports Create subfolders Schedule reports Create data worksheets, dashboards, and reports
Reports	Schedule Reports	View dashboards and reports Create subfolders Schedule reports
Reports	View Reports	View dashboards and reports Create subfolders

User Management Permissions

The following table lists the permissions needed to manage users.

Function	Permissions	Allows users to...
User Management	View Users	View the list of all active and inactive users
User Management	Create Users	View users
		Assign roles to users
		Assign users to groups
User Management	Activate /Deactivate Users	View users
		Change the status of a user that you manage
User Management	Change User Password	View users
		Change the password of a user that you manage
User Management	Change User Email	View users
		Change the email associated with a user
User Management	Assign Roles to Users	View users
		Assign roles that you currently have to users that you manage

Function	Permissions	Allows users to...
User Management	Assign Users to Groups	View users View account groups Add and remove users from account groups that you currently manage Assign users who are members of account groups that you manage to any other account group
User Management	Manage Groups	View account groups Create account groups *You are automatically added to the account groups that you create.* Delete account groups that you currently manage Add and remove managers for account groups that you currently manage Add and remove users from account groups that you currently manage Assign users who are members of account groups that you manage to any other account group
User Management	Manage Roles	View roles Create roles *You are automatically added to the account groups that you create.* Add and remove users from roles that you have Add and remove any permission assigned to you from roles that you currently have Delete roles that you currently have

ArcSight Permissions

The following table lists the permissions available when you deploy an ArcSight capability such as Log Management and Compliance.

Function	Permission	Allows users to…	Available with…
ArcMC	ArcMC System Admin *Not available in a SaaS environment*	Perform System Admin functions	Fusion
ArcMC	ArcMC Operation Admin *Not available in a SaaS environment*	Perform all Operations functions, but does not have access to System Admin	Fusion
ArcMC	ArcMC System Viewer *Not available in a SaaS environment*	Read only access to System Admin functions	Fusion
ArcMC	ArcMC Operation Viewer *Not available in a SaaS environment*	Read only access to Operations functions	Fusion
Dashboards	Share a dashboard	With the Manage Roles permission, share the current dashboard with any role Without the Manage Roles permission, share the current dashboard with any of the roles associated with the user's role	Fusion
Licensing and Usage	Manage Contract	Create and edit an MSSP profile Import, update, view, and delete an MSSP contract	an MSSP license
Licensing and Usage	Access EPS Usage	Export an EPS Usage Report	an MSSP license
Searches	Execute Search	Execute searches using fieldsets, custom ranges dates, and search operators	Fusion
Searches	Export Search Results	Export the search results in csv format	Fusion
Searches	Never Expire Search Results	Configure searches to never expire	Fusion
Searches	Manage Scheduled Searches	Create and manage scheduled searches	Fusion

Function	Permission	Allows users to...	Available with...
Searches	Import / Export Search Queries	Import and export search queries	Log Management and Compliance
Searches	Import / Export Search Criteria	Import and export search criteria	Log Management and Compliance
Searches	Perform Event Integrity Check	Run an Event Integrity Check and view the results	Log Management and Compliance
Searches	Manage Outlier Models and Scoring	Create and delete Outliers models Build and pause the scoring processes	Log Management and Compliance
Searches	Manage Lookup Lists	Add, configure, view, and delete lookup lists	Fusion
Searches	Manage Fieldsets	Create, edit, and delete fieldsets	Fusion
Searches	Manage Search Queries/Criteria	Create, clone, edit, delete,and view all previously saved search queries and search criteria View and clone all out-of-the-box search queries	Fusion
Searches	Logger Data Migration *Not available in a SaaS environment*	Execute a data migration from Logger into the ArcSight Database	Fusion
Operations Management	Access Database Monitoring-Overview	View high-level, summary information about the workload and health of the database	Capabilities that require the ArcSight Database
Operations Management	Access Database Monitoring-Details *In a SaaS environment, available only to the System Operations Administrator*	View details about the health of the individual components of the distributed database system	Capabilities that require the ArcSight Database
Operations Management	Manage Storage Groups	Create and manage storage groups	Fusion
Operations Management	Manage Kafka *Not available in a SaaS environment*	Access Kafka Manager for Transformation Hub	Transformation Hub

Default Roles

Fusion provides several default roles. If you have the *Manage Roles* permission, you can change the permissions of any role assigned to your account except for those of the *System Admin*. You can also create additional roles that reflect your organization's needs.

Some permissions are available with specific functionality, such as Reports or Search, or when you have a particular license, such as for Log Management and Compliance.

> As of the Fusion 1.4 release, some roles are no longer default roles. However, Fusion continues to display them if you deployed your environment before the roles were deprecated. For example, *ArcMC User*, *Guest, User*, and *Report User* are no longer default roles.

Default Role	Permissions
System Admin *Not available to customers in a SaaS environment*	All permissions
Admin	All **Dashboard** permissions All **Licensing and Usage** permissions All **Reports** permissions All **Searches** permissions All **User Management** permissions *Access Database Monitoring-Overview*

Default Role	Permissions
Analyst	All **Dashboard** permissions
	Execute Search
	Manage Fieldsets
	Manage Search Queries/Criteria
	Schedule Reports
	View Reports
System Operations Administrator *Not available to customers in a SaaS environment*	*Access Database Monitoring-Overview*
	Access Database Monitoring-Details
	All **Dashboard** permissions
	All **ArcMC** permissions
	Manage Kafka
SCIM Integration Read-Only *Not available to customers in a SaaS environment. Not created in other environments*	View Users

Create a Role with Permissions

You can group multiple permissions into a role and assign the relevant role to your users. A user must have at least one role.

You can assign only the permissions and roles that you have yourself.

1. Click **ADMIN** > **Roles and Permissions** > **Create Role**.
2. In the field in the upper left corner, specify a name for the role.
3. Press **Enter**.
4. Select the **permissions** that you want to apply to the new role.
5. To add users to the role, complete the following steps:
   a. Select the **USERS** tab.
   b. Select **Assign role to users**.
   c. Choose the users you want to add to the role.
   d. **Save** your changes.

View Details of a Role

When you view the details of a role, you can also modify the role's settings and permissions.

1. Click **ADMIN** > **Roles and Permissions** > *role_name*.

2. (Optional) Modify the role in one of the following ways:

   - Change the set of permissions

     You can assign only the permissions and roles that you have yourself.

   - Add or remove users

   - Delete the role

Change Permissions for the Role

You can only assign permissions that you have yourself.

1. While viewing a role, select **Permissions**.
2. In the **Permissions** tab, select the permissions that you want to add or remove.

   You might need to scroll the page to see the full set of available permissions.

Add or Remove Users for the Role

You can add or remove multiple users in a role.

1. While viewing a role, select **Users**.

2. In the **Users** tab, select **Assign role to users**.

3. Select the users that you want to assign to or remove from the role.

   You can also add or remove roles for a specific user.

Delete the Role

While viewing a role, select **Remove role from users**.

You can delete any role except the *System Admin* role.

Using Visuals and Reports to Analyze Data

Your environment must include a capability that uses the reports.

The **Reports Portal** allows you to browse and filter your datasets and to visualize results in the Portal's reports and dashboards. Rapidly discover meaningful trends and associations that yield actionable intelligence. The built-in Admin reports enable a report administrator to track use of the Portal.

If your product provides built-in reports and dashboards, you usually can find them in the *Standard Content* directory of the Portal's repository. Depending on your assigned permissions, you can view, schedule, design, or manage reports and dashboards. You add custom reports and dashboards by collecting and filtering data from your connected sources. The Reports Portal supports the ability to drill down into specific elements for thorough data reviews.

Accessing Reports and Dashboards in the Reports Portal

Your environment must include a capability that uses the Reports Portal. Also, you must have one of the Reports permissions *to use this feature.*

Select **Reports** > **Portal**.

When you view the dashboards and reports, be aware that they are not persistent. Once you leave a report or dashboard, you must regenerate the view when you return to the page. If you choose to open a report in a new browser tab, you can leave that tab open to keep the dashboard or report active while you look at other dashboards or reports.

Many out-of-the-box reports and dashboards contain pre-built queries. When you run a report or view a dashboard, it might prompt you to provide values for the run-time parameters. Reports also prompt for the start and end time of the data search.

View a Dashboard

When you open a dashboard, it automatically retrieves data from the last two hours. However, you can modify the time range as needed.

1. Select **Reports** > **Portal** > **Repository** > **Standard Content**.

2. Expand the desired category, then select the dashboard that you want to view.

3. (Optional) To change the time range for the report, modify the start or end time parameters.

   When you change the time range, the dashboard refreshes the data.

View a Report

When you open a report, you must define the time range for the data that you want to view.

1. Select **Reports** > **Portal** > **Repository** > **Standard Content**.

2. Expand the desired category, then select the report that you want to view.

3. To specify the time range, complete the following steps:

   a. To activate the Calendar, point your cursor at the position of the **Calendar** icon to the right of the time selection box.

   b. Select the **Calendar** icon.

   c. Enter the **Start Time** for the report.

   d. Enter the **End Time** for the report.

4. Select **Submit**.

   The report will execute and display when it is complete.

5. (Optional) To email the report when it completes, select **Schedule**, then define the delivery options.

Specify Default Dashboards for the Reports Portal

The Reports feature allows you to specify the default dashboards that display when you enter the Reports Portal. You can choose from any of the content available within the Reports Repository. Alternatively, if you have the *Design Reports* permission, you can create dashboards that you or others might want to include in their default dashboard.

For example, in the Reports Portal, you might want a ready access to dashboards that you use regularly. So you add the OWASP Missing Security Patches Overview and Foundation Denial of Service Activity dashboards.

1. Select **Reports** > **Portal** > **Portal Dashboards**.

2. Specify a name for your default dashboard.

3. (Optional) Enter a description for your dashboard portal.

4. Select one of the available dashboards.

   You can specify only one dashboard at this time. However, once you are in the Reports Portal, you can add more dashboards. Each dashboard appears as a tab in the page.

5. (Conditional) To create a dashboard, select **Compose Dashboard**.

6. Click OK.

7. (Conditional) If you chose to create a dashboard, continue adding the items that you want to include. For additional instructions, select (?).

Designing Reports for Data Analysis

*You must have the **Report Admin** or **Design Reports** permission to use this feature.*

Select **REPORTS** > **Designer**.

Report **Designer** provides a wizard that allows you to create new reports using the bundled Standard Content data worksheets. You can design elements, change their attributes, and control all aspects of element presentation and layout. The Designer saves all attributes and related information in a template file in XML format. The Designer also supports visually building queries against multiple types of data sources and specifying data grouping, summarization and element data binding.

The Designer offers you the same functionality as an API, but makes most tasks, such as report layout, much simpler. You can also use the Designer to attach scripts to embed business logic into the report.

Scheduling Report Generation

*You must have the **Report Admin** or **Schedule Reports** permission to use this feature.*

Select **REPORTS** > **Scheduler**.

The Reports **Scheduler** enables you to schedule and manage batch report generation. You can create one or more scheduled tasks for which you specify a time condition, reports to be generated, and delivery mechanism of the generated output.

The Reports feature can output the reports in formats such as PDF and Excel. The Scheduler can send the reports in email, save to disk or an archive, or print them.

Adding and Removing Reports Content

*You must have the **Report Admin** permission to use this feature.*

Select **REPORTS** > **Content**.

The Reports **Content** enables administrators to modify the reports and dashboards in the following ways:

- Add and remove content, also known as assets, for the reports and dashboards using the **Import Assets** and **Export Assets** feature.

- Connect to data sources using the **Add Data Source** feature. Using this feature, you can gather content from specific sources to supply reports and dashboards.

Import and Export Content

This capability is not available in a SaaS environment.

Use the **Import Assets** and **Export Assets** options to manage the reports and dashboard available to your users. You can move assets from one server environment to another. For example, you might want to move a set of reports from a test server to a production server.

> **Note:** If Reporting generates errors when you attempt to export assets, you should reduce the number of assets that you export concurrently.
>
> Alternatively, you might need to increase the RAM for the Reporting node. For more information about sizing your environment for the workload, see the Technical Requirements for the ArcSight Platform. However, in a SaaS environment you will not be able to adjust the RAM for the Reporting node.

> **Note:** You cannot export content from the **My Reports** folder.

Supported Data Sources

This capability is not available in a SaaS environment.

You can incorporate data from the following sources:

Text/Excel Directory

Connects to a specified file (text or Excel) or file location.

To access and upload this file type, you must create a new folder for your files in the /var/lib/inetsoft/ path on the reporting server. You might need assistance from your Server Admin.

REST JSON

Connects to a REST (Representational State Transfer) data source containing JSON (JavaScript Object Notation)-formatted data.

REST XML

Connects to a REST data source containing XML-formatted data.

JDBC

Connects to a relational database using Java Database Connectivity.

This source supports commercial and open source databases such as Oracle, SQL Server, DB2, Sybase, Informix, MySQL, PostgreSQL, and MS Access. Be sure to download the latest driver (https://www.inetsoft.com/support/drivers.jsp).

Elasticsearch REST

Connects to an open source search engine.
The process for adding this type of data source is the same as for adding an Elasticsearch data source.

R

Connects to an R database containing R language sources.

Best Practices for the Report Designer and Dashboard Designer

When using the Reports Portal, follow these best practices to improve your work flow for creating reports and dashboards.

Use Search Results to Create a Dashboard or Report

Each completed search has a unique **Search Results ID**, which represents a link to the temporary table containing the search results. You can copy that ID, then build a report or dashboard around the search results.

- Build a Report Using Search Results
- Build a Dashboard Using Search Results
- Convert the Search Fields to Human-Readable Values

Build a Report Using Search Results

You can build a report around results of a previously run search by leveraging the Search Results ID.

1. When viewing the Events table for a search, select the **Copy** icon in the table's header.

   This icon contains the **Search Results ID**.

2. Select **Reports** > **Report Designer**.

3. Select **Create** > **Report**.

4. In the **Select a data source** field, paste the Search Results ID that you copied.

   The retention period of the temporary table in the database is 30 days.

5. (Optional) Convert the fields in the temporary table to human-readable values.

6. Continue creating the report.

Build a Dashboard Using Search Results

You can build a dashboard around results of a previously run search by leveraging the Search Results ID.

1. When viewing an Events table, select the **Copy** icon in the table's header.

   This icon contains the **Search Results ID**.

2. Select **Reports** > **Dashboard Designer**.

3. Select **Create** > **New Dashboard**.

4. From the visual composer, select **Data Source** > **Database** > **TABLE** > **Default_secops_recon**.

5. Select the ID of the search that you previously copied.

   The retention period of the temporary table in the database is 30 days.

6. Select **Open wizard** or **OK**.

7. (Optional) Convert the fields in the temporary table to human-readable values.

8. Continue creating the dashboard where the Search Results ID is the data source.

Convert the Search Fields to Human-Readable Values

The ArcSight Database uses a temporary table to store content associated with a Search Results ID. Because the names of the fields in the table represent the coding-style name, you might want convert the terms to more user-friendly values.

To change the field names, your report or dashboard must use a Data Worksheet.

1. Select **Reports** > **Dashboard Designer**.

2. Open the dashboard or report that you want to modify.

3. From the upper-right corner, select the **Data** icon.

4. Open the worksheet.

5. In the lower pane, select the **Formula Editor** icon. The tool-tip for this icon says "Create Expression."

6. Select **SQL**.

7. In the Expression pane of the Formula Editor, add the following strings:

```
Time: to_timestamp(field['normalizedEventTime']/1000)
IP:   v6_ntoa(field['sourceAddressBin'])
MAC:  mac_btoa(field['sourceMacAddressBin'])
```

8. Select **OK**.

9. In the lower pane of the worksheet, select the **Change Data Mode** icon.

10. Select **Live Event** data.

11. Hide the binary (original) fields.

12. **Export** or **Save** the dashboard or report as needed.

Use Data Models to Build a Worksheet

Select **Reports** > **Reports Designer** > **Report type** > **Data Source** > **Database**.

Data models are logical models of the events table in the database that allow for an extra level of abstraction where you can perform varied transformations. You can use the final data model as the final table when creating a data worksheet. By default, the system has two data models:

Basic Data Model

Contains fewer fields from the events table. Use this model for an easier understanding or for simple reports that require less fields.

Event View

Contains the entire events table.

You can also create, edit, and delete your own Data Models. For more information, see "Create a Data Model" in the Help in the Reports Portal. Make sure to add only the fields you that need and create the filters from there. Some of the fields in the data model are non-human readable. You should parse them to ensure that they are readable in the report.

Use Data Worksheets to Build a Dashboard or Report

Data worksheets define the base for the reports and dashboards. Using data worksheets allows you to freely manipulate different data origins and generate a final set of results that can be used for reports and dashboards.

1. Select **Reports** > **Dashboard Designer** or **Report Designer**.

2. From the upper-right corner, select the **Data** icon.

3. From the right corner, select the **New Data Worksheet** icon.

4. To start the worksheet, complete one of the following actions:

    4a (Conditional) To browse for a data source, select **Database Query**, then **OK**.

    4b (Conditional) To import a data file, select **Upload File**, then **OK**.

    4c (Conditional) To open a new worksheet then choose the data source, select **Mashup Data**, then **OK**.

    4d (Conditional) To open a new worksheet, select **Cancel**.

5. Drag and drop the fields, tables, or queries that you want to include in the dashboard or report.

    Alternatively, you can create tables, then link them using unions or joins.

    > Using joins to show correlations between data sources like CSV files and event charts might cause slow performance depending on the size of the files. For larger data sources, see Use Pre-Populated Search Results.

6. (Conditional) To refine the design, select one of the following options from the Preview pane.

    For example, you can sort and reorder the columns or change the data mode.

    > Be sure to hide or remove fields that you don't need for your dashboard or report.

7. To save your changes, complete the following steps:

    7a Select **Save** or **Save As**.

    7b Specify the folder where you want to save the worksheet.

    Do not specify the **Standard Content** folder, which is reserved for the built-in reports and dashboards.

When you create a custom report, do not base that report on any existing **Standard Content** DataWorksheet. Instead, create a new DataWorksheet or use the DataSource/table selection options.

Create a Simple Dashboard

When creating a simple dashboard, Reports prompts you to select the data source. When you open the Dashboard Visual Composer, a window displays where you can choose the data source for the Dashboard. Follow the prompts or close the window to continue to the main editor of the Dashboard.

From the Dashboard editor, you can create Tables and Charts in the canvas. From there, you can also convert to measure some fields that can provide numeric values and can be used in a chart. You can also convert to dimension the fields that can provide a string value.

First, use the system to create and save a data worksheet as the basis for your dashboard. Use one of the following to create a simple dashboard.

- Use the Dashboard Wizard
- Use the Dashboard Editor

Use the Dashboard Wizard

If you select the wizard, the Dashboard Designer displays the Wizard section of the Dashboard. From here, you can create the first component of the Dashboard.

1. Select **Reports** > **Dashboard Designer** > **Crosstab Wizard**.
2. Select the **data worksheet** of your preference as a data source, and then click **Next**.
3. Select **Open Wizard**.
4. Select the fields to use in your dashboard.
5. (Conditional) Select the dashboard style:

   Crosstab

   Groups the dashboard by row and column headers and displays the summary data at the intersections

   Table

   Groups the dashboard and summarizes it or displays it in tabular layout

   Chart

   Creates multiple charts using multiple fields

   Full Editor

   Allows granular control view of your updates, such as format, color, and shape

6. Once the editing is complete, set the position of the element in the dashboard canvas.

7. View the dashboard, and then select **Continue**.

8. Once the dashboard has been successfully edited, select **Finish**.

9. Click **Save as** to save your dashboard.

Use the Dashboard Editor

Using the Dashboard Designer, you can edit the elements and freely set their position in the Dashboard. The Dashboard Designer displays the Wizard section of the Dashboard.

1. Select **Reports** > **Dashboard Designer** > **Crosstab Wizard**.

2. Click **Cancel** to open the dashboard editor.

3. Select the **data worksheet** of your preference as a data source, and then click Next.

4. Add the elements available from the left.

5. Update the dashboard using the Dashboard composer.

   You can create, add, and edit multiple elements.

6. Click **Save** to save your dashboard in a **Custom Content** folder.

Create a Parabox Chart

Parabox charts, also called Parallel Coordinates, allow you to visualize connections between multiple entities and identify significance based on the size of target elements. This allows you to quickly discover threats and respond to suspicious activity in your environment.

1. Select the **Create New** icon.

2. Choose the data source.

   Make sure the data source includes the fields you want to represent on the parabox chart.

3. When the visual composer opens, drag the chart from the left panel.

4. Select the **Edit** icon.

5. Select **Full Editor**.

6. Drag the fields you want for the Y axis to the box labeled **Y**.

7. Right click on the Y axis.

8. Select **Hide Title**.

9. Right click on the chart.

10. Select **Properties** .

11. Select **Script**.

12. In the **Script** section, paste the code below.

```
if(this.data.length>1){
        //new EGraph() creates a new instance of a EGraph object and
represents the graph definition
        graph = new EGraph();
        //Create a new Parabox element
        var elem = new ParaboxElement();
        //Obtain fields and ordering from binding dialog instead of
having to declare them
        var y =  bindingInfo.yFields;
        var scales = [];
        for(var j=0; j < y.length; j++) {
                var z = new CategoricalScale(y[j]);
                //hide line
                z.getAxisSpec().setLineVisible(false);
                //add categoricalScale to array
                scales.push(z);
                //Add the required fields to the parabox.
                elem.addParaboxField(y[j])
                }
```

```
                var coord = new ParaboxCoord(scales[0],scales[1]);
                coord.setScales(scales);
                //set vertical spacing between data points; default is 10,
requires build 143798+
                //coord.setSpacing(15);
                //Set the Font and Color of the column labels
                var labelColor = coord.getAxisLabelScale().getAxisSpec
().getTextSpec();
                labelColor.setFont(java.awt.Font('Roboto',java.awt.Font.PLAIN,
10));
                labelColor.setColor(java.awt.Color(0xFFFFFF));
                //Sets the graph coordinates to the Parabox Coordinates defined
above
                graph.setCoordinate(coord);
                //Create new Text Specifications to set the text format of the
points
                var pointColorFont = new TextSpec();
                pointColorFont.setFont(java.awt.Font
('Roboto',java.awt.Font.PLAIN, 11));
                pointColorFont.setColor(java.awt.Color(0xFFFFFF));
                elem.setTextSpec(pointColorFont);
                //set shape and color of data points
                var sframe = new StaticShapeFrame();
                var cframe = new StaticColorFrame();
                cframe.setColor(java.awt.Color(0x0073E7));
                sframe.setShape(GShape.FILLED_CIRCLE);
                elem.setShapeFrame(sframe);
                elem.setColorFrame(cframe);
                //set connecting line color
                var cframeLine = new StaticColorFrame();
                cframeLine.setColor(java.awt.Color(0x44495B));
                elem.setLineColorFrame(cframeLine);
                //Add the parabox element to the graph
                graph.addElement(elem);
                }
```

13. Select **Apply**.

Create a Simple Scheduled Report

You can create a report that runs on your chosen schedule. In the report, define conditions that trigger tasks and actions you want to run.

1. Select **Reports** > **Scheduler**.

2. In the lower left corner of the screen, select **New Task**.

3. For **Name**, enter a name of the task.

4. To set the conditions for your report, complete the following steps:

   a. Select the **Condition** tab.

   b. (Conditional) To specify the timezone that the report uses, perform one of the following actions:

      - To use the timezone where the server is installed, select **Show Server Time Zone**

      - To use your timezone, deselect **Show Server Time Zone**

   c. (Conditional) To run the task at specific intervals, configure the frequency.

      For example, to run a report every Monday afternoon, specify the following settings:

      - Select **Time Range**, then **Afternoon**

      - For **Every**, enter 1

      - Select Monday

   d. (Conditional) To run the tasks in sequence, select **Chained**, then specify the first task.

   e. Select **OK** to save the scheduled task.

5. To specify the report associated with the scheduled tasks, complete the following steps:

   a. Select the **Action** tab.

   b. For **Report**, click **Select** then navigate to the report that you want to schedule.

   c. To email the report results, select **Deliver to Emails** then configure the email content and destination addresses.

   d. To set the time range in which the report retrieves data, complete one of the following actions:

      - Select **Add**, and then specify the time values.

      - Select **Creation Parameters**, then choose the dates from the calendar option.

   e. Select **OK** to save your changes.

Create a Simple Report

First, create and save a data worksheet. For additional details on how to create a data worksheet, see Using Data Worksheets to Build a Dashboard or Report.

Use the one of the following wizards to create a simple report.

- Use the Crosstab Wizard
- Use the Table Wizard
- Use the Chart Wizard
- Guidelines for Report Usage

Use the Crosstab Wizard

From the Reports Designer menu, use the Crosstab Wizard to create a report that displays data in a pivot table where the data is grouped by row and column headers, and the summary data is displayed at the intersections.

1. Select **Reports** > **Report Designer** > **Crosstab Wizard**.
2. Select the **data worksheet** of your preference as a data source, and then click **Next**.
3. Define the **row and column groups** (vertical and horizontal columns), and then click **Next**.
    - For **Row groups**, select the row headers.
    - For **Column groups**, select the column headers.
4. (Conditional) Define the **summary columns** that will display as summarized fields.
5. (Conditional) **Filter the conditions** that will define the original data.

    After the design statement is filled, the options for insert, modify, and clear will be enabled.
6. (Conditional) For **table style**, use the default option.
7. To complete the editing, click **Finish Editing**.

Use the Table Wizard

From the Reports Designer menu, use the Table Wizard to create a report that displays data in tabular layout or grouped and summarized.

1. Select **Reports** > **Report Designer** > **Table Wizard**.

2. Select the **data worksheet** of your preference as a data source.

3. Select the columns to display in the report from the select **detail columns**.

4. Define the groups to display as **column headers**.

5. (Conditional) Define the **summary columns** that will display as summarized fields.

6. (Conditional) Filter the conditions to define the original data. Once the design statement is filled, the control options are enabled.

7. (Conditional) Retain the default **table style** for better formatting results.

8. (Conditional) Rank the groups to display as top or bottom groups.

Use the Chart Wizard

From the Reports Designer menu, use the Chart Wizard to create a chart-based report.

1. Select **Reports** > **Report Designer** > **Chart Wizard**.

2. Select the **data worksheet** of your preference as a data source.

3. By default, the auto option is selected. Use the **chart style** to style your report.

4. (Conditional) If required, select one of the following 2D and 3D images chart styles.

   Your chart options include bar, line, area, point, pie, donut, radar, stock, candle, box plot, waterfall, pareto, map, treemap, and marimeko charts.

5. Define the **X Axis** that to display as columns.

6. Define the **Y Axis** to display as columns.

7. Define the visual properties (color, shape, size, text) of the columns by using the visual binding.

8. (Conditional) Filter the conditions to define the original data. Once the design statement is filled, the control options are enabled.

   (Conditional) Rank the groups to display as top or bottom groups.

9. (Conditional) Additional steps might be required depending on the chart style selected:

   Geographic binding

   Use if you select **Map Style** for your report. Choose different aspects about the map report that will be generated.

   Tree dimensions

   Use if you select **Treemap**, **Sunburst**, **Circle Packing**, or **Icicle** for your report. Select

the fields

the report will use for the Tree Mapping.

Marimekko category

Use if you select Marimekko Style for your report. Select the field for the Marimekko Category Dimension.

Guidelines for Report Usage

- Create as many data models as needed but only include the fields that you need for your report. For simple reports, use the Basic Data Model instead of the event view.

- To convert non-human readable fields in the data model, parse them before adding them to the report.

- You can create filters from the data model or the report itself. It is recommended to set the filters from the data model so these can be saved in the data base.

- Check the meta data box for a faster pre-visualization of the report. Take into consideration that no real data is displayed with this option.

- Export the results in CSV format for faster results.

- When needed, copy the built-in dashboards and use them as templates for other creations.

Improve the Performance of Dashboards and Reports

You can improve the performance of your reports, dashboards, and worksheets by following these best practices.

- Use Raw Database Fields instead of Defined Functions Fields
- Use `normalizedEventTime` Instead of the Time Field
- Use the Integer Variant Instead of the String Variant
- Display Host Names Instead of IP Addresses on Charts
- Use `startswith` or `endswith` Instead of `contains` to Create Conditions
- Put the Expensive Conditions at the Top of your Worksheets
- Put the Most Expensive Conditions Towards the Top of your Building Blocks Hierarchy
- Put Parameters at the Top of your Worksheets
- Use the Flyover Option for Dashboards with Multiple Charts and Tables

Use Raw Database Fields Instead of Defined Functions Fields

Where possible, use raw database fields over defined function fields to speed up the search process by limiting the number of events searched.

f icons represent the defined function fields. 🮱 icons represent the raw database fields.

For example, use: `[Events.deviceAddressBin][is not][null]` instead of `[Device Address][is not][null]`.

Use Normalized Event Time Instead of the Time Field

Use `normalizeEventTime` instead of the **Time** field from the logical model. Because the **Time** field requires extra calculations, whereas `normalizeEventTime` is a raw field, your query will run more quickly. For more information, see Use Raw Database Fields instead of Defined Functions Fields.

Use the Integer Variant Instead of the String Variant

When data can be represented in a string format or an integer format, use the integer format of the data field because strings are defined functions and integers are raw database fields. For more information, see Use Raw Database Fields instead of Defined Functions Fields.

For example, use: `[Events.agentSeverity][is][one of][3,4]` instead of `[Events.Agent Severity String][is][one of][High,Very-High]`.

Display Host Names Instead of IP Addresses on Charts

Where possible, use host names because host names are represented as raw fields in the database. IP addresses represented as database function `v6_ntoa(Events1_0.destinationAddressBin) AS Target Address` will be calculated for every selected event.

Use 'startswith' or 'endswith' instead of 'contains'

To create conditions, use `startswith` or `endswith` instead of `contains`, when possible. This narrows your search, and your queries will process more quickly.

For example: `[Events.categoryTechnique][is][starting with][/Traffic Anomaly]`.

Put the Most Expensive Conditions at the Bottom of Your Worksheets

When you have conditions that take up a lot of operation space, put them at the bottom of your worksheet. This will limit the quantity of events that the expensive condition must search, and thus speed up your query.

Put the Most Expensive Conditions Towards the Top of Your Building Blocks Hierarchy

For advanced users.

When you need to display expensive conditions on your dashboard, you can move it up in the hierarchy. This improves your dashboard's performance because those conditions will not run against every event.

> **Note:** This does not work for reports, only dashboards.

> **Caution:** If you have a wrapper function and move it up the in the hierarchy, you need to define the wrapper function multiple times. For example, if you have multiple charts showing the same field from different angles (one chart is Top Target IPs, and another chart for relationships between attacker IPs and Target IPs) you will need to define the wrapper function twice.

Put Parameters at the Top of your Worksheets

If you create a dashboard or report with parameters (for example, data for a specific host), arrange the worksheet conditions so that the parameter goes before the complex conditions. This will limit the events that the complex conditions search and speed up your query.

For example:

[Events.destinationHostNamelowerCase][is][contains][$(hostname(equal_or_like))]

[and]

[Events.Time][is][between][$start_time.$(end_time)]

Use the Flyover Option for Dashboards with Multiple Charts and Tables

When you have a dashboard with a table, consider using the flyover option instead of using the charts and tables on one screen. If you designed the dashboard with the flyover option, the dashboard will show one chart. When you click on a specific target, a flyover table will show the information for this specific target. Right-click the dashboard you want to use the flyover option.

1. Select **Properties**.

2. Under **Flyover**, click the box or boxes for the additional chart or table you want to appear, and then click **OK**.

Additionally, when you again pause the mouse over the specific target, the flyover information appears more quickly because it is drawn from the cache.

When to Write your own SQL Query

If you want to run a very specific query, and you do not want to use the logical model, you can write your own SQL query.

1. From New Data Worksheet, select Database Query.

2. Add fields.

3. Write your SQL query, save it, and use it like any other worksheet. For more information, see Use Search Results to Create a Dashboard or Report.

Searching for Events

The **Search** feature enables you to look for and investigate events that meet specified criteria so you can detect anomalies that point to security threats. You can view the results in tabular and timeline formats, as well as view the raw event data. Each search consists of specifying query input, search result fields, and the criteria for which you want to search events.

Queries are case sensitive. The query input determines the search type (full text, natural language, or contextual). As you specify the criteria for a search query, Search suggests items and operators based on a schema data dictionary. You can also choose from predefined search queries.

Understanding Search

The application ingests log data, migrated from ArcSight Logger and SmartConnectors, that has been routed through Transformation Hub and events from ArcSight Enterprise Security Manager. Each entry in a log is referred to as an **event**. The application accepts events from Transformation Hub and organizes them to maximize search and storage efficiency.

The **Search** feature enables you to look for and investigate events that meet specified criteria so you can detect anomalies that point to security threats. You enter a search query, the criteria (such as a time window) over which to search, and the fields from the Unified Event Schema. You can use one of the three timestamps the database stores for each event for your time window.

Search displays results in an Events timeline as a histogram chart, which shows the number of events returned over event occurrence time. The Events table shows events returned by search. The table displays columns of **fields**, each representing a particular categoryof data, such as an IP address or the port where the event originated. When you select an event, you can view its list of field-value pairs in the Event Inspector panel. For ongoing or regular searches, you can save queries, queries plus specific criteria, and search results. You can also schedule searches to run on a regular basis.

For the query's time range, you can choose a fixed start and end date, where you cannot refresh data, or a predefined date range. For example, for the last **30 minutes** predefined search, you receive updates upon re-executing the search based on the most recent 30 minutes. Alternatively, you could specify dynamic dates, such as **Midnight on the first day of the current month**.

After initiating a search, you can pause, resume, and cancel the process as needed. A progress bar shows you the percent of retrieved data.

> Because search results consume space, the system maintains a threshold for the total number of executed searches that it can store. Stored content includes saved search results, completed runs of session searches, and completed runs of scheduled searches. The system displays a notification when the threshold is passed. If this occurs, you cannot run a search and scheduled searches cannot run until some previously executed searches have been deleted.

Understand Session versus Saved Searches

Select **Search**.

As you initiate searches, Search automatically preserves your activity in case you must navigate to another search or to a different feature in the ArcSight Platform. Search temporarily maintains these **session searches** in tabs until you close the search tabs, exit your browser, or log out. The Home tab lists all your current session searches. Therefore, if you close the search tabs or lose the search tabs by logging out, you can open them again from the Home tab.

- "Saved Searches" below
- "Expired Searches" below

Saved Searches

For long-term use of a search, you must save the query, criteria, or results. You can review and manage your **saved searches** at any time:

- Saved queries
- Saved criteria
- Saved search results

Expired Searches

Session and saved searches usually **expire** after a specified amount of time. When they expire, the system deletes all information about the search. If this occurs when you have the search open in a tab, you will receive a notification that the search has expired and be instructed to close the tab.

Because session searches are considered short-term or temporary searches, the default expiration time is after 24 hours of inactivity. Saved searches expire after seven days by default. You can reset the expiration time by running the search again or by modifying the query or criteria. However, if the search has already expired, you cannot reset the expiration clock.

You can also override the default expiration time by changing the **Search expires in** setting for a particular session or saved search.

Understand the Search Progress Indicators

As **Search** retrieves data, it displays a **progress bar** to show its status, including the percent of data received. Rather than attempting to read all data at once, Search gathers data in chunks of time. The progress bar shows the time range from which the results are currently being retrieved.

You can **pause the search** and restart as needed.

NOTE: When performing a search with two or more identical queries the number of events returned for the second search will correspond to the next chunk of data. If you pause then resume the search, the first search will be moved to the next chunk as well, maintaining the same number of events retrieved. The identical queries can contain either one of the built-in queries or a custom query.

Understand the System Searches

Search includes the following out-of-the-box **system searches** that contain a query plus specific criteria. All of these system searches are set in Normalized Event Time. For more information about how to use these queries and criteria, see "Load a Saved Search" on page 146.

> **Note**: You can also search queries by using # and the query name. For example, #Configuration changes or #DGA Events. Additionally, you can run criteria searches as queries using the same method. Additionally, there is a list of reserved words that must be enclosed in quotes (" ") to ensure the system correctly parses the query.

Category	Name	Use as	Description
Application Monitoring	Windows New Service Created	Query	Lists events indicating new windows services were created from the following event sources: • Microsoft-Windows-Security-Auditing:4697 • Service Control Manager: 7045
Configuration Monitoring	Configuration changes	Query	Lists configuration changes based on ArcSight categorization.
Entity Monitoring	Failed logins	Query	Lists events indicating failed login activity based on ArcSight categorization.
	Failed Login Events	Criteria	Lists failed login activity events based on ArcSight categorization for the last 30 minutes by default.
	Failed logins for $username	Query	Lists events indicating failed login activity based on ArcSight categorization for a specific user. The user should be specified before running the search.
	Windows account creation	Query	Lists events indicating new windows accounts created based on the following event sources: • Microsoft-Windows-Security-Auditing:4720 • Security:624
Event Monitoring	ESM Correlation Events	Query	Lists ESM correlation events.
Malware Monitoring	Malicious code activity	Query	Lists events indicating malicious code activity based on ArcSight categorization.

Category	Name	Use as	Description
MITRE Monitoring	MITRE ATT&CK Events	Criteria	Lists correlation events reported from Arcsight ESM content package: https://marketplace.microfocus.com/cyberres/content/esm-default-content. These events are forwarded to the ArcSight Database using ArcSight Forwarding connector, or any other flex connector which reports this information, using the following mapping: deviceCustomString6Label='MITRE ID' Where deviceCustomString6 contains the actual MITRE ATT&CK technique.
Network Monitoring	DGA Events	Criteria	Lists DGA-related events based on Microsoft Trace Log.
	DNS Events	Query	Lists DNS-related events.
	DoS Events	Criteria	Lists events indicating denial of service based on ArcSight categorization.
	Firewall drop	Query	Lists Drop Firewall events based on Arcsight categorization for a specific IP address. The IP address should be provided at runtime.
	Firewall drop for $ip		Lists Drop Firewall events based on Arcsight categorization.
	Firewall Events	Criteria	Lists Firewall events based on ArcSight categorization.
	Proxy Events	Criteria	Lists Proxy events based on ArcSight categorization.
	SSH authentication	Query	Lists events indicating SSH Authentication events based on ArcSight categorization.
	VPN connections	Query	Lists events indicating VPN connections based on ArcSight Categorization.
Vulnerability Monitoring	Vulnerabilities Events	Criteria	Lists events indicating vulnerabilities based on ArcSight categorization and Vulnerability Scanner events.

Understand Search Queries

You must have the **Manage Search Queries** *permission.*

A **search query** is a set of conditions used to select events when you run a search. For example, you can enter a very simple term to match such as "login" or an IP address. Alternatively, you can specify a complex query to match events that include multiple IP addresses and reference a lookup list. In the search query, you can enter the alias, or abbreviated term, for a field name rather than entering the full name. You can also use the **presentable field names**, such as Agent Address.

Your query input determines the search type: full text, natural language, or contextual. As you specify the fields and values for the query, Search suggests search items and operators based on a schema data dictionary.

Search provides default queries, labeled as *system*. However, you can save your own queries, which you can load into another search. You have the option to clone, modify, or remove a saved query at any time.

Understand the Query Syntax

Depending on the type of search you create, your query must meet the requirements listed in the following table. Search treats a comma (,) between search items and values as an **OR** operator. Additionally, there is a list of reserved words that must be enclosed in quotes (" ") to ensure the system correctly parses the query.

If you do not get the search results you expect, you might need to restate the query. For example:

- If the query is written with spaces, only the first word is shown in the results. A better way to write the query statement is to use explicit phrasing without any spaces.
- Queries that filter specific "id" field values (for example, id = "123456789" or id != "123456789") will not return correct results. Create the query without using "id" fields.

By default, search operations are case-sensitive to support faster performance. However, you can instruct the database to support case-insensitive searches. For SaaS deployments, talk to your SaaS Admin about changing the database.

When you construct a query, you can include operators, such as eval and lookup, for more robust searches.

> You cannot use multiple operators, such as *NN* and *XX*, in the same query.

- "General Syntax Rules" below
-

General Syntax Rules

Type	Full-text	Field-based	Hashtag (predefined)
Case sensitivity	Case-sensitive	Case-sensitive	Case-insensitive
Exact Match	Keyword treated as keyword*. Example: /Execute matches: /Execute, /Execute/Start, /Execute/Response,/Execute/Query	Enclose value in double quotes. Example: `Category Behavior ="/Execute"`	n/a

Nesting, including parenthetical clauses, such as (a OR b) AND c	Allowed Use boolean operators to connect and nest keywords.	Allowed Use boolean operators to connect and nest keywords.	Allowed Use boolean operators to connect and nest keywords.
Implicit Operators	When you enter two values separated by a space, this is treated as an implicit AND condition. Example: `ssh fail`	The AND/OR treatment depends on the operator used in the search. For example, destinationAddress = 1.1.1.1, 2.2.2.2 is equivalent to destinationAddress = 1.1.1.1 or destinationAddress = 2.2.2.2 , while the query destinationAddress != 1.1.1.1, 2.2.2.2 is equivalent to destinationAddress != 1.1.1.1 and destinationAddress != 2.2.2.2	n/a
List Operations	n/a	Performs an inner join or a left join against a custom list. *Syntax for an Inner Join:*`source address in list CustomListName_ CustomColumn Name` *Syntax for a Left Join:*`source address not in list CustomListName_CustomColumnName`	n/a
Time Format (when searching for events that occurred at a particular time)	No specific format The query needs to contain the exact timestamp string. Example: `"10:34:35"`	YYYY-MM-DD YYYY-MM-DD HH:mm YYYY-MM-DD HH:mm:ss.fff To narrow the time range, use the following operators: • in between (><) • greater than (>) • less than (<)	n/a

Special Characters: \ * ' "	Use the backslash (\) as an escape character.	Use the backslash (\) as an escape character.	n/a
Wildcard	Can appear anywhere in the value. Examples: *log log* lo*g* Searches for ablog, blog, long, etc.	Can appear anywhere in the field. Examples: name=*log Searches for ablog, blog, etc. in name field name="*log" name=*log Both search for *log	n/a
Escape a Wildcard Character	Can search for * by escaping the character. Example: log*	Can search for * by escaping the character. Example: log*	n/a

Implicit Operators

Implicit operators form the basic building blocks for query construction. Use them along with other operators and functions to create robust search queries.

To build queries, use the following general operators:

Operator	Alternative Operator	Examples
AND		#Firewall drop and sourceAddress equals 10.0.112.9 sourceAddress equals 10.0.112.9 and destinationAddress = 10.0.116.148
OR		fail OR ssh destinationAddress = 10.0.111.5 OR destinationAddress=10.0.116.148 destinationAddress =10.0.111.5, 10.0.116.48
not equal	<> !=	destinationPort not equal 21

equals	= == is equal to equal	name equals INVALID password device vendor equals CISCO
greater than	> is greater	bytes In greater than 100
less than	< is less is lower less	bytes out less than 1000
greater equal than	>= gte greater equal	End Time greater equal than 2017-07-25 End Time greater equal than 2017-07-25 09:07 End Time greater equal than 2017-07-25 09:07:43 End Time greater equal than 2017-07-25 09:31:22.685
less equal than	<= lte less equal	Base Event Count less equal than or equal 50
starts with	startwith	message starts with FIN
does not start with		name does not start with FIN
ends with	endswith	message ends with out
does not end with		message does not end with out
contains	contain like has substring	name contains TCP
does not contain	does not have	name does not contain TCP
in list	match in list of	device vendor equals CISCO and source address in list customListName_ customColumnName device vendor equals CISCO and source address in list badGuyIpList_badGuyIp
not in list	not match not in list of	source address not in list customListName_ customColumnName source address not in list badGuyIpList_ badGuyIp
in subnet	n/a	source address in subnet 10.0.0.0/8
not in subnet	n/a	source address not in subnet 10.0.0.0/8

Understand the Types of Search Queries

Search supports the following types of search queries:

- "Full Text Search" below
- "Field-based Search" below
- "Hashtag (predefined searches)" below

Full Text Search

Searches across all fields using a 'contains' operation to determine if the value is found.

Syntax	Example
<value>	ssh

Field-based Search

Searches based on the field and operator designation to determine if the value is found in the specified field.

Your search can reference fields with the Unified Schema to either retrieve the field in results, apply a filter criteria or create a user defined expression. The **Unified Schema** defines a consistent event model that can be used across all of ArcSight family of products.

Syntax	Example
<key> <operator> <value>	sourceAddress = 10.0.111.5

Hashtag (predefined searches)

The Search feature includes several predefined queries out-of-the-box. In the query field, enter a hashtag, and then select the criteria to use. In addition to these predefined searches, you can use the session searches and save searches in the input field using a hashtag prefix.

To ensure the system correctly parses your query, if your search entity name includes one of the reserved words listed before, you should surround the query name with quotes (" ") in order to avoid ambiguity in the query statement.

This predefined query...	Description
#Configuration Changes	Lists configuration changes based on ArcSight categorization.
#DGA Events	Lists DGA-related events based on Microsoft Trace Log.
#DNS Events	Lists DNS-related events.
#DoS Events	Lists events indicating denial of service based on ArcSight categorization.
#ESM Correlation Events	Lists ESM correlation events.
#Failed Logins	Lists events indicating failed login activity based on ArcSight categorization.
#Failed Logins For User $Username	Lists events indicating failed login activity based on ArcSight categorization for a specific user. The user should be specified before running the search.
#Firewall Drop	Lists Drop Firewall events based on Arcsight categorization for a specific IP address. The IP address should be provided at runtime.
#Firewall Drop For $Ip	Lists Drop Firewall events based on Arcsight categorization.
#Firewall Events	Lists Firewall events based on ArcSight categorization.
#Malicious Code Activity	Lists events indicating malicious code activity based on ArcSight categorization.
#MITRE ATT&CK Events	Lists correlation events reported from Arcsight ESM content package: https://marketplace.microfocus.com/cyberres/content/esm-default-content. These events are forwarded to the ArcSight Database using ArcSight Forwarding connector, or any other flex connector which reports this information, using the following mapping: deviceCustomString6Label='MITRE ID' Where deviceCustomString6 contains the actual MITRE ATT&CK technique.
#Proxy Events	Lists Proxy events based on ArcSight categorization.
#SSH Authentication	Lists events indicating SSH Authentication events based on ArcSight categorization.
#VPN Connections	Lists events indicating VPN connections based on ArcSight Categorization.

#Vulnerabilities Events	Lists events indicating vulnerabilities based on ArcSight categorization and Vulnerability Scanner events.
#Windows Account Creation	Lists events indicating new windows accounts created based on the following event sources: • Microsoft-Windows-Security-Auditing:4720 • Security:624
#Windows New Service Created	Lists events indicating new windows services were created from the following event sources: • Microsoft-Windows-Security-Auditing:4697 • Service Control Manager: 7045

Use Reserved Words in a Query

To ensure the system correctly parses your query, if your search entity name includes one of the reserved words listed before, you should surround the query name with quotes (" ") in order to avoid ambiguity in the query statement.

For example, if your query name is: "System warnings and errors" use the following notation: #"System warnings and errors."

Reserved Words for Queries		
and	as	between
by	category	connecting to
contain	contains	custom float
distinct	does not contain	does not end with
does not start and end with	domain	ends with
endswith	equal	equals
filter	for	greater
greater equal	greater equal than	greater than
gte	has	has substring
hostname	ibt	id
in between	in cidr block	in list
in list of	in subnet	is
is between	is equal	is equal to
is greater	is greater or equal than	is greater than
is greater than or equal to	is larger	is larger than
is less	is less equal	is less or equal than
is less than	is less than or equal to	is lower
is lower than	is not	is not between
is not equal	is not equal to	ip
ip6	label	less
less equal	less equal than	less than
like	lte	mac

Reserved Words for Queries		
match	nibt	not
not between	not equal	not equals
not in between	not in cidr block	not in list
not in subnet	not match	not within subnet
or	path	pipe
port	span	starts and ends with
starts with	startswith	timestamp
username	uri	url
where	wheresql	within subnet
withinsubnet		

Include a Storage Group's Filter in the Search Query

Search allows you to include a storage group in a query. For example, you have a storage group called *Firewall Events* that has the following query: `categoryDeviceGroup='/Firewall' or categoryDeviceGroup='/IDS'`. Rather than entering that query again in Search, specify the following for your Search query: `storageGroup=Firewall Events`.

> ⚠ **IMPORTANT**: For best results, specify the storage group at the beginning of the Search query.

Use GlobalEventID in a Query

To help you identify an event that might be seen by multiple ArcSight components, the connectors assign the event a unique 64-bit ID. To include a GEID in your search query, enter globalEventID. You can view the GEID of the event in the Event Details.

Syntax	Example
global event id=<value>	global event id= 2864991913017849867

For events to have a GEID, use ArcSight Management Center to configure connectors to include the ID. For more information in:

- Non-SaaS environments, see "Generator ID Manager" in the *Administrator's Guide for the ArcSight Platform*
- SmartConnectors, see "Unique Generator ID" in the *ArcSight SmartConnector Installation Guide*.

Specify an Alias for a Field

In the search query, you can enter the alias, or abbreviated term, for a field name rather than entering the full name. For the fields shown in the following table, you can also use the **presentable field names**, such as Agent Address. Search suggests presentable names.

Field	Aliases
agentAddress	agt
	agent ip
agentHostName	ahost
agentId	aid
agentMacAddress	amac
	agent mac
agentReceiptTime	art
agentTimeZone	atz
agentTranslatedAddress	agent translated ip
agentType	at
agentVersion	av
applicatonProtocol	app
	protocol
baseEventCount	cnt
bytesIn	in
bytesOut	out
categoryBehavior	behavior
categoryDeviceGroup	device group
categoryObject	object
categorySignificance	significance
categoryTechnique	technique

Field	Aliases
destinationAddress	dst
	destination ip
	destinationip
	dst ip
	dest ip
	target ip
	targetip
	target
destinationHostName	dhost
	destination name
destinationMacAddress	dmac
	destination mac
destinationNtDomain	dntdom
destinationPort	dpt
	destination port
	dstport
	dest port
	targetport
	target port
destinationProcessId	dpid
destinationProcessName	dproc
destinationTranslatedAddress	destination translated ip
destinationuserId	duid
destinationUserName	duser
	dst user
	dest user
	destination user
	dst usr
destinationUserPrivileges	dpriv
deviceAction	act

Field	Aliases
deviceAddress	dvc
	deviceaddr
	deviceip
	device ip
deviceCustomFloatingPoint*n*	cfp*n*
Valid values for *n* are integers between 1 and 4	For example: cfp1
For example: deviceCustomFloatingPoint1	
deviceCustomFloatingPoint*n*Label	cfp*n*Label
Valid values for *n* are integers between 1 and 4	For example: cfp1Label
For example: deviceCustomFloatingPoint1Label	
deviceCustomIPv6Address*n*	c6a*n*
Valid values for *n* are integers between 1 and 4	device custom ipv6 *n*
For example: deviceCustomIPv6Address2	For example: c6a2
deviceCustomIPv6Address*n*Label	c6a*n*Label
Valid values for *n* are integers between 1 and 4	For example: c6a2Label
For example: deviceCustomIPv6Address2Label	
deviceCustomNumber*n*	cn*n*
Valid values for *n* are integers between 1 and 3	For example: cn3
For example, deviceCustomNumber3	
deviceCustomNumber*n*Label	cn*n*Label
Valid values for *n* are integers between 1 and 6	For example: cn6Label
For example: deviceCustomNumber6Label	
deviceCustomString*n*	Cs*n*
Valid values for *n* are integers between 1 and 6	For example: Cs5
For example: deviceCustomString5	
deviceEventCategory	cat
deviceHostName	dvchost
deviceMacAddress	dvcmac
	device mac
deviceProcessId	dvcpid
deviceReceiptTime	rt
deviceTimeZone	dtz

Field	Aliases
deviceTranslatedAddress	device translated ip
endTime	end
eventOutcome	outcome
fileNme	fname
fileSize	fsize
message	msg
requestUrl	request URL
sourceAddress	src source ip sourceip src ip
sourceHostName	shost
sourceMacAddress	smac source mac
sourceNtDomain	sntdomain
sourcePort	spt srcport src port
sourceProcessId	spid
sourceProcessName	sproc
sourceTranslatedAddress	source translated ip
sourceUserId	suid
sourceuserName	suser src user source user src usr
sourceUserPrivileges	spriv
startTime	start
transportProtocol	proto

Specify a Group of Fields

Search enables you to quickly select fields that have common groupings. In the query, you can specify a **group alias** that displays all field associated with the group. The following table provides some common group aliases.

Group Alias	Includes a list of these fields...
category	All category fields
custom float	All custom float fields
domain	All domain fields
hostname	All hostname fields
id	All ID fields
ip	All IP address fields
ip6	All IPv6 address fields
label	All label fields
mac	All MAC address fields
path	All path fields
port	All port fields
timestamp or time	All time fields (device receipt time, agent receipt time)
uri	All URI fields
url	All URL fields
username or user	All user fields

Specify IP Addresses and Subnets

Your query can include IPv4, IPv6, and MAC addresses. Search stores IPv4, IPv6, and MAC addresses in a format that provides search flexibility and enables you to perform the following actions:

Compare IP addresses for optimum performance

For example, `Agent Address > 192.10.11.12`.

Specify a range of IP addresses

For example, you can enter the following types of queries:

- `Agent Address in between 192.2.13.1 and 192.2.13.11`
- `Source Address greater equal than 192.10.11.12`
- `Destination Address less than 192.112.98.33`

Specify a range of IP addresses

For example, you can enter the following types of queries:

- `Agent Address in between 192.2.13.1 and 192.2.13.11`
- `Source Address greater equal than 192.10.11.12`
- `Destination Address less than 192.112.98.33`

Use abbreviated input search notation

You can enter the following types of queries:

- To specify IP addresses in the subnet starting with a particular value:
  `Agent Address in subnet 192.*`
- To specify an IPv4 address in a subnet that uses CIDR notation. The first eight bits are the network part of the address, leaving the last 24 bits for specific host addresses.
  `Agent Address in subnet 192.0.0.0/8`
- To specify an agent address in a subnet that uses CIDR notation. The first 24 bits are the network part of the address, leaving the last 40 bits for specific host addresses.
  `Agent Address in subnet 2001:0db8:0000:0000:0000:ff00:0042:8329/24`

Search stores MAC addresses in their original format.

To enter an IP or MAC address in a search query:

Enter the MAC addresses in the following formats:

- aa:aa:aa:aa:aa:aa
- aa-aa-aa-aa-aa-aa

The following table lists the query format and examples for the type of IP address.

Type of address	Format in a query...	Examples
IPv4	a.b.c.d	a.* a.b.* a.b.c.* a.b.c.d/8
IPv6	Full form	2001:0db8:0000:0000:0000:ff00:0042:8329
	Canonical form without leading zeroes in each group	2001:db8:0:0:0:ff00:42:8329
	Canonical form without consecutive sections of zeroes	2001:db8::ff00:42:8329
IPv6 in a subnet	Include CIDR notation	2001:0db8:0000:0000:0000:ff00:0042:8329 2001:0db8:0000:0000:0000:ff00:0042:8329/24 2001:db8::/32 NOTE: For the 2001:db8::/32 format, you can omit part of the IPv6 address, depending on the subnet that you are querying.
MAC	a:b:c:d:e:f a-b-c-d-e-f	94:18:82:6D:63:74 94-18-82-6D-63-74

Use an Operator in the Query

Create powerful queries with search operators and functions. You can also select several out-of-the-box system searches that contain a query plus specific criteria. Operators, such as eval, can be chained together to create complex queries.

> Search operators and functions must be entered in all lower case letters when they are used in queries.

> Do not use a raw event field as part of a query.

Use Cases for Search Operators

The following are just a few examples of the flexibility and power of search operators.

-
-

You may need to adjust a query to work with your own fieldsets.

For more information about working with operator chaining see and .

General Search Operator Use Cases

I want to see where possible brute force password guessing is happening.

Additional Information: To determine this, I want to see the top 10 devices that are responsible for the most number of failed logins.
Operator used: top

```
#FailedLogin | top 10 deviceEventClassId
```

I want to know the hourly amount of data transfer on MyWebserver.

Operators used: chart, sum, by, span

```
sourcehostname = MyWebserver.com | chart sum(bytesIn), sum(bytesOut) by
deviceVendor, deviceProduct span 1h
```

I want to see a sum of events, grouped by hostname and day.

Operator used: chart
Aggregate function: sum (This summarizes the values passed as an input, grouped by the "by" clause.)
Time bucket: 1h (Events are grouped in time increments of one hour.)

```
| chart sum(baseEvents) by hostName span 1h
```

I want to determine all account lockouts, grouped by user name.

Operators used: wheresql, top

```
(deviceVendor="Microsoft" and deviceProduct="Microsoft Windows") or
deviceProduct="Unix" | wheresql  deviceEventClassId in
["Security:539","Security:644","arcsight:66:0","Microsoft-Windows-
Security-Auditing:4740","Microsoft-Windows-Security-Auditing:6279"] and
destinationUserName is not null |top destinationUserName
```

Operator Chaining Use Cases

I want to identify the rare occurrences of Firewall events.

Additional Information: I want to determine this from 3 specific fields' data (device vendor, category device group, and name).
Operators used: rename, rare (bottom)

```
#Firewall Events | rename deviceVendor as DV | rename category device
group as CDG | rare DV , name , CDG
```

I want to isolate vulnerabilities.

Additional Information: I will base this on data from 3 significant fields (device vendor, category technique, and device group), then determine the most common occurrences found in those categories.
Operators used: rename, rare (bottom)

```
#Vulnerabilities | rename deviceVendor as DV | rename category technique
as CT | rename category device group as CDG | rare DV , name , CDG , CT
```

I want to apply filters to a set of fields and then to extract the top-50 most common occurrences of those events.

Operators used: where, top

```
source address is not null | where Bytes In >= 3000 | where Category
Outcome = /Success | top 50 source address , Category Outcome
```

I want to determine the top insecure processes on devices in my company.

Operators used: top, rename

```
destinationProcessName in ["telnetd", "ftpd", "pop3", "rsh" ,
"imapd","rexec"] | top destinationProcessName  | rename
destinationProcessName as  "Process"
```

```
deviceVendor = ArcSight | rename sourceUserName USER | top USER
```

Show me all configuration changes by product.

Operators used: top, rename

```
categoryBehavior = "/Modify/Configuration" and categoryOutcome =
"/Success" | top deviceProduct | rename deviceProduct_count_2 as "Changes"
| rename deviceProduct as "Product"
```

I want to apply filters to a set of fields and then to extract the top-50 most common occurrences of those events.

Operators used: where, top

```
source address is not null | where Bytes In >= 3000 | where Category
Outcome = /Success | top 50 source address , Category Outcome
```

Chaining Search Operators

Construct a complex query statement by chaining together multiple search operators into a single query instead of implementing separate queries. This powerful capability lets you perform robust, real-world searches while providing the flexibility to customize searches for specific scenarios. You can save these searches to reuse them in future updates.

Operator chaining is a process by which the search takes a set of results from one operation and uses these results as input for the next operation. Chaining a series of operations equips you with the options needed to "slice and dice" data to extract and analyze it on a highly granular level. Operator chaining works with all the pipeline operators (rename, eval, where/filter, wheresql, top, bottom/rare and chart/stats).

During operator chaining, fieldsets become more restricted as more operators are added to the query, especially with eval and aggregation operators. For example:

```
severity!=null | top severity | stats avg (Count_1) by severity
```

For information about operator chaining workflows, see "Use Cases for Search Operators" on page 86.

Syntax Recommendations

Use the following syntax recommendations to ensure operator chained searches succeed.

- To use the fields from a **lookup list** table with the search operators, make a **join** with one of the lookup fields using the **"in list"** operator. You also should add the lookup list fields to the current fieldset. For example:

  Add a lookup list with name as **Customer** then add its field, which will be used with search operator (e.g. **Customer_Vendor**) to the current fieldset.

  ```
  Source Address in list Customer_Address | wheresql Customer_Vendor =
  'Microsoft'
  ```

- **Alias/New** field name cannot be an existing field name or a synonym of an existing field name. Also, an alias field name cannot be an existing group name or reserved word.

  In this example, "destination hostname" and its synonyms "dhost" and "destination name" cannot be used as aliases.

- **Alias/New** field name should not have spaces (like test 1), otherwise it will cause conflicts. These are examples for acceptable alias/new field names:

```
name is not null| eval test1 = concat(name, "_test") | eval test_2 =
upper(test1)
```

```
name is not null | where name not equals ARCSIGHT | chart count (distinct
name) as Dcount by name
```

- The following is an example of how to use a generated field with **eval** in another operator:

```
| eval test = upper ( name ) | where test != "ARCSIGHT"
```

- **Count_<number>** cannot be used as an alias for a field name.
- The **wheresql** operator is case sensitive. Just like all other operators, the wheresql name must be stated in all lower case letters.
- Do not create new field name with spaces if these new fields will be used later with the **wheresql** operator. The where condition of wheresql operator will not recognize new field with spaces that were created by previous operators. In addition to wheresql, this is also applicable for the **eval** and **chart and stats** operators. Here are two **invalid** examples:

```
| rename name as new name | wheresql new name = 'TCP'
```

```
destination port is not null | eval convert name = upper ( Name ) |
wheresql convert name = 'MSTYPE'
```

- You can use the **where** operator to filter dynamic fields, for example:

```
| top 5 Name | where Count_1 > 1000
```

- More filters can be added to a search through the Fields Summary feature. Click on **Fields Summary** and select a field and a value for that field. The new filter will be appended at the end of the query in use as a |**where** clause.
- Multiple **aggregate functions** can drastically modify drastically the fieldset that is available for the next pipe operator. for example:

```
name is not null | char count (Name) by Device Vendor span 1h | chart
count (Name) by Name
```

The second chart pipe cannot access to span operations because the NET, DRT, dBRT are not available for this chaining level.

Same scenario applies to the **top** operator:

```
name is not null | top Name | char count(count_Name_1) by Name span 1h
```

chart/stats

The **chart/stats** operators display search results for specified fields as fields in the Search Results table.

- "Syntax" below
- "Aggregation Functions" on the next page
- "The Span Function" on page 93
- "How Do I Use This?" on page 94

Syntax

```
...| chart count by field1, field2, field3 ... [span [time_field]=time_
bucket]
...| chart {{sum | avg | min | max | } (field)}+ by field1, field2, field3
...[span [time_field]= time_bucket]
...| chart {function (field)} as new_field_name by field [span [time_field]
```

For simple syntax examples, see below.

where:

- *field*, *field1*, *field2* are the names of event fields used in system queries.
- *time_bucket* is the bucket size (in any combination) used for grouping events. Use **d** for day, **h** for hour, **m** for minute, and **s** for seconds. For example, *2h*, *5d*, *1m*.
- *function* is one of the following: **count**, **sum**, **avg** , **min**, **max**, **latest** or **earliest**..
- *new_field_name* is the name you want to assign to the field (field of data) in which the function's results are displayed. For example, *Total*.
- All chart/stats commands accept only one field in the input. For example, | stats count (device vendor) by...
- The input field must contain a field that exists in the database.
- If multiple fields are specified, separate the field names using commas without any spaces.
- The function input field must contain numeric values for chart/stats sub-operators that are mathematical operations (sum, avg, min, max).
- The mathematical operators **avg** and **mean** are equivalent.

"by" Statements

- The chart/stats operators and eval, require a "by" statement. For example: eval | chart sum(AgentSeverity) by Destination HostName

- Specify a field name after the "by" statement. For example: "by deviceVendor"

- Fields in a "by" statement should be separated by comma. For example: ... | chart count (Name) by deviceEventCategory , name

Simple Syntax Examples

- Grouping events by a field and counts how many names each group has.

  ```
  ...| chart count(name) by Destination Hostname
  ```

- Sum of Bytes In for every group of events. These events (rows) are grouped by the field Destination Hostname.

  ```
  ...| chart count(Bytes In) by Destination Hostname
  ```

- The events (rows) are grouped using the default time field selected, distributing the events in groups of 1 hour. The events that matches the same time bucket and the same Agent Severity will be organized in the same group. Then returns how many names contains every group.

  ```
  ...| chart sum(name) by Agent Severity span = 1h
  ```

Aggregation Functions

> Aggregation functions only work on numeric fields. The specified fields must contain numeric values. If a field you specify is of the wrong data type, you will receive an error message like the following: "java.lang.NumberFormatException".

- You can include multiple functions in the same chart/stats command. When doing so, separate each function with a comma, as shown:

- ```
  ... | chart count (Name) by Destination Hostname, sum
  (deviceCustomNumber3) by deviceEventClassId
  ```

- When you include multiple functions, the search results table displays one field per function.

- You can use the "**as new_field_name**" clause to name any field resulting from the aggregation functions, as shown:

  ```
  ...| chart sum(deviceCustomNumber3) as TotalStorage, avg
  (deviceCustomNumber3) as AverageStorage by deviceCustomNumber3
  ```

- Instances of "**new_field_name**" should be changed by the alias name of the aggregation result.

Aliases that contain special characters have the following syntax restrictions:

Special Characters	Restrictions	Examples
+, *, - , &, =, !, | , <, >	Enclose in single/double quotes when they are reused and the search runs as expected.	| rename file path as 'FP+DEV' | chart count ('FP+DEV') by 'FP+DEV'
@, #, $, %, {}, [], ?, _, ^, /, ~, ., \	These do not need to be enclosed in single/double quotes when they are reused and the search runs as expected.	| rename file path as 'FP$DEV' \| chart count (FP$DEV) by FP$DEV
\	When a backslash is uses as part of an escape clause, add an additional backslash \. It does not need to be enclosed in single/double quotes when they are reused and the search runs as expected. The outcome field name should show only one backslash.	| rename file path as 'FP\\DEV' | chart count (FP\\DEV) by FP\\DEV

The Span Function

In addition to grouping events defined by **eval** operators, the **span** function groups events by a time field (such as EventTime or deviceReceiptTime) and a time bucket.

- The span function can only accept three fields (Normalized Event Time, Device Receipt Time, and Database Receipt Time) before the equals sign, for example:
  - |chart sum(Agent Severity) by Destination hostName span Normalized Receipt Time = 1h
  - |chart sum(Agent Severity) by Destination hostName span = 1h
  - |chart sum(Agent Severity) by Destination hostName span 1h
- The span operator is not allowed after an aggregation operator.
- The span operator must use an equal sign or a supported field name. For example: span Normalized Event Time = 1h
- The span operator only accepts timestamp fields. For example: span Normalized Event Time = 1h.

- A span's time bucket must use one of the following types of values: *number*d, *number*h, *number*m, and *number*s

- By default, the chart/stats command displays the first 10 unique values. If the span function creates more than 10 unique groups, not all of them will be displayed.

- When span is included in a query, search results are grouped by the specified time bucket. For example, if span=5m, the search results will contain one row for each 5-minute span. If there are no events within a specific 5-minute span, that row will be empty.

- The span function assumes a **24-hour day, all year long**. If span=1d or 24h, on the day of the daylight savings time change, the event time indicated by the span_eventTime field in the search results will be different from the previous day by one hour. On the day when there are 23 hours in a day (in March), the span bucket will still include events from the last 24 hours. Similarly, on the day when there are 25 hours in the day (in November), the span bucket will include events from the last 24 hours.

How Do I Use This?

Aggregation Function Examples

- Use the default chart setting (Column Chart) to specify multiple fields. In this example, a count of unique groups of deviceEventCategory and name fields is displayed and plotted.

```
... | chart count(Name) by deviceEventCategory name
```

- Simple query using 1 (min) aggregation function and one group by field (min)

```
... | [chart | stats] aggregation_function (field_name) [as alias_name]
by field_name
```

```
 chart count (Name) by Destination Hostname
```

- Query using 1 to *N* aggregation functions, the result and 1 to *N* group by field (min)

```
... | [chart | stats] [aggregation_function (field_name), ...
aggregation_function (field_name)] [as alias_name] by field_name
```

```
 chart count (Name), sum(Agent severity) by Destination Hostname
```

- Query using 1 to *N* aggregation functions, the result and 1 to *N* group by field (min)

```
... | [chart | stats] [ aggregation_function (field_name), ...
aggregation_function (field_name)] [as alias_name] by field_name [,
```

```
...field_name]
```

```
chart count (Name), sum(Agent severity) by Destination Hostname, Name
```

- Adding alias in every aggregation function

```
... | [chart | stats] [ aggregation_function (field_name), ...
aggregation_function (field_name)] [as alias_name] by field_name [,
...field_name ]
```

```
chart count (Name) as alias_name1, sum(Agent severity) as alias_name2 by
Destination Hostname, Name
```

Span Function Examples

- If time stamp is specified for span input, it does not use parenthesis: The correct query would be:

```
...| chart count(Name) by deviceEventCategory span deviceReceiptTime = 5m
```

- Destination Hostname is the time field and one hour is the time bucket:

```
 chart count (Name), sum(Agent severity) by Destination Hostname, Name
span 1h
```

- deviceReceiptTime is the time field and 5m (5 minutes) is the time bucket:

```
...| chart count(Name) by deviceEventCategory span (deviceReceiptTime) =
5m
```

- Span is used to organize events by time frame:

```
... | [chart | stats] [ aggregation_function (field_name), ...
aggregation_function (field_name)] [as alias_name] by field_name [, ...
field_name ] span = n[s|m|h|d]|m|h|d]
```

- If a time field is not specified for the span function , EventTime is used as the default. For example, the following query uses EventTime by default:

```
...| chart count(Name) by deviceEventCategory span = 5m
```

Grouping with span is useful in situations when you want to find out the number of occurrences in a specific time span.

- If you want to find out the total number of incoming bytes every 5 minutes on a device, you can specify a span of 5m. This example assumes that deviceCustomNumber1 field provides the incoming bytes information for these events.

```
...| chart sum(deviceCustomNumber1) by hostName span 5m
```

- You want to see a sum of events by hostName in one week of events, listed by day. When a **span** field is specified in conjunction with an **event** field, the unique sets of all those fields are used for grouping.

```
...| chart sum (baseEvents) by hostName span = 1d
```

- The following example uses deviceCustomNumber and deviceAddress in conjunction with span to find out the number of events (using deviceCustomNumber3) from a specific source (using deviceAddress) in one (1) hour:

```
...| chart sum(deviceCustomNumber3) by deviceAddress span=1h
```

- When span is included in a query, search results are grouped by the specified time bucket. For example, if span=5m, the search results will contain one row for each 5-minute span. If there are no events within a specific 5-minute span, that row will be empty.

- The span function assumes a **24-hour day, all year long**. If span=1d or 24h, on the day of the daylight savings time change, the event time indicated by the *span_eventTime field* in the search results will be different from the previous day by one hour. On the day when there are 23 hours in a day (in March), the span bucket will still include events from the last 24 hours. Similarly, on the day when there are 25 hours in the day (in November), the span bucket will include events from the last 24 hours.

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

rename

Use the pipeline operator **rename** to assign a new name to a portion of the search query.

- "Syntax" below
- "How Do I Use This?" on the next page

Syntax

```
... | rename source_name as new_source_name
```

where

- *source_name* represents the field that you want to rename.
- *new_source_name* represents the new name that you want to apply to the field.

Aliases that contain special characters have the following syntax restrictions:

Special Characters	Restrictions	Examples
@, #, +, ?, /, ^, [], {}, _ , *, $, % , ., ~ , \	Do need to be enclosed in single/double quotes when they are reused and the search works as expected.	| rename source address as 'source@' | rename source@ as SA
&, !, - , = , <, >, |	Need to be enclosed in single/double quotes when they are reused and the search works as expected.	| rename source address as 'source&' | rename 'source&' as SA
\	When a backslash is uses as part of an escape clause, add an additional backslash \. It does not need to be enclosed in single/double quotes when reused and the search runs as expected. The outcome field name should show only one backslash.	| rename source address as 'source\\' | rename source\\ as SA

How Do I Use This?

- Assign a new address name to an existing source address.

```
... | rename source_address as SourceAddressABC
```

For more information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

top and bottom

The **top** and **bottom** operators list the search results of the most common values for the specified field. The resulting values are listed in tabular format from the highest count value to the lowest.

The fields can be event fields, available in the application menu. If multiple fields are specified, you need to separate the field names with white space or a comma.

- "top" below
- "bottom" below
- "Syntax" below
- "Parameters" on page 101
- "How Do I Use This?" on page 101

top

The **top** operator provides the most common values for the specified field(s). The values are listed from the highest count value to the lowest.

bottom

The **bottom** operator provides the least common values for the specified field(s). The values are listed from the lowest count value to the highest. The **rare** operator can be used as an alias to **bottom**.

Syntax

```
…| top [N] field1 [,field2, field3]
```

where:

- [ ] indicates optional input you may enter.
- Italicized characters indicate where a user can enter *custom* field information.

- Only *N* is optional. *N* limits the matches to the top *n* values for the specified fields. One (1) field is required, but you can specify a maximum of five (5) integers, separated by commands.
- If you do not specify *N*, the default value is 500.
- If included, *N* should be between one (1) and the search results limit.
- The operator performs a standard count (*) to retrieve the number of events.
- No search operator other than "where" can be used in a query after the top/bottom operator is used.
- Queries that use the top/bottom search operator along with fields that begin with "Device" may fail completely or partially. To avoid this behavior, select the field from the drop-down options that are available as you enter the query. This also applies to fields that are not editable.

Aliases that contain special characters have the following syntax restrictions:

Special Characters	Restrictions	Examples
*, -	Do not need to be enclosed in single/double quotes when they are reused and the search works as expected.	Destination port <> null | rename Destination Port as 'D*P' | rename Source Port as 'S*P' | top 10 D*P , S*P
@, #, +, ?, /, ^, [], {}, _ , ~, \, ., $, %, .	Do not need to be enclosed in single/double quotes when they are reused and the search works as expected.	Destination port <> null | rename Destination Port as 'D#P' | top D#P
&, ! , = , <, > , +, |	Need to be enclosed in single/double quotes when they are reused and the search works as expected.	Destination port <> null | rename Destination Port as 'D=P' | top 'D=P'
\	When a backslash is uses as part of an escape clause, add an additional backslash \. It does not need to be enclosed in single/double quotes when reused and the search runs as expected. The outcome field name should show only one backslash.	Destination port <> null | rename Destination Port as 'D\\P' | top D\\P

Case 1: alias name with special character * , - does not need to be enclosed in single/double quotes when reused and search works as expected. Ex: Destination port <> null | rename Destination Port as 'D*P' | rename Source Port as 'S*P' | top 10 D*P , S*P

Case 2: alias name with special character @, #, +, ?, /, ^, [], {}, _ , ~, \, . , $, % does not need to be enclosed in single/double quotes when reused and search works as expected. Ex: Destination port <> null | rename Destination Port as 'D#P' | top D#P

Case 3: alias name with special character &, ! , = , <, > , +, | needs to be enclosed in single/double quotes when reused and search works as expected. Ex: Destination port <> null | rename Destination Port as 'D=P' | top 'D=P'

Parameters

The parameters are *N* and a list of comma-separated fields.

For the **top** operator, when multiple fields are specified, the count of unique sets for all of the fields is listed from the highest to lowest count. For the **bottom** operator, the fields are listed from the lowest to the highest count.

How Do I Use This?

The top operator is used to limit the matches to the top *N* values for the specified fields. Likewise, the bottom operator is used to limit the matches to the bottom *N* values for the specified fields. The default count number is 500 unless you specify a value for *N*. Here are a few examples:

- You want to limit your results to the 1,000 most common event categories.

  ```
  …| top 1000 deviceEventCategory
  ```

- You want to limit your search for the top 5 event categories.

  ```
  …| top 5 categories
  ```

- You want to see all products from a specific vendor that are sending the least number of events.

  ```
  deviceVendor = Vendor| bottom 10 deviceProduct
  ```

- See the "rare" user action in the organization happening using the HTTPS protocol.

  ```
  protocol=https | rare requestuseragent
  ```

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

where

The **where** operator displays events that match criteria specified in a "where" expression. Where expressions act as filters to return only those results that fulfill a particular condition. In fact, **filter** is a synonym of the operator **where**. Results for where expressions are binary, satisfying either true or false.

- "Syntax" below
- "How Do I Use This?" on the next page

Syntax

```
... | where <expression>
```

where:

- The *where* operator represents the filter you want to use us on a field.
- The *expression* field represents a valid field-based query expression. Arithmetic expressions or functions are not supported.
- For *where any … contains* queries, all fields are executed, but only results for alpha (letter) IDs are displayed. For example, results for the ID "HostName" display, but results for the ID CEID-3631 will not display, even though the field is executed.
- You can specify multiple field conditions in one query expression by using the listed operators between them. The conditions can also be nested. For example:

  ```
  (name="John Doe" OR name="Jane Smith")AND message!="success"
  ```

Aliases that contain special characters have the following syntax restrictions:

Special Characters	Restrictions	Examples
@, #, +, ?, /, ^, [], {}, _ , *, ., ~ , \, $, %, .	Do need to be enclosed in single/double quotes when they are reused and the search works as expected.	| rename source address as 'source@' | where source@ <> null
&, !, - , =, < , >, |	Need to be enclosed in single/double quotes when they are reused and the search works as expected.	.| rename source address as 'source&' | where 'source&' <> null

Special Characters	Restrictions	Examples
\	When a backslash is uses as part of an escape clause, add an additional backslash \. It does not need to be enclosed in single/double quotes when reused and the search runs as expected. The outcome field name should show only one backslash.	| rename source address as 'source\\' | where source\\ <> null

How Do I Use This?

```
... | where eventId is NULL
```

```
... | where eventId=10006093313 OR deviceVersion CONTAINS "4.0.6.4924.1"
```

```
... | where eventId >=10005985569 OR categories= "/Agent/Started"
```

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

wheresql

The **wheresql** operator supports all of the Eval functions listed in this guide. The Database also supports many SQL functions for a "where expression that you might want to use. Nested queries are allowed, but you may use only one dataset for the search. If using more than two expressions, use parentheses to nest expression clauses.

- "Syntax" below
- "Parameters" below
- "How Do I Use This?" below

Syntax

```
…|wheresql boolean_expression
```

Wheresql expressions are binary, satisfying either true or false. You must construct your queries with syntax supported by the ArcSight Database.

Parameters

You can include the following parameters:

- AND (&&)
- OR (||)
- NOT (|)
- LIKE

How Do I Use This?

- You want to construct filter your search results to display numerical data between 10 and 50.

```
…|wheresql bytesOut between 10 and 50
```

- Match the company name of a device vendor.

```
… | wheresql regexp_ilike(deviceVendor, 'Company_Name')
```

- Further refine the search to select a device vendor from a specified table used in another search.

```
… | wheresql deviceVendor in (select deviceVendor from tableX)
```

Syntax Recommendations

Use the following syntax to ensure searches and schedule searches that use the "wheresql" condition succeed:

- As with all search operators, the operator name must be in all lower case letters.
- Fields must have valid names as listed in the ArcSight Database.
- Enclose **string** values in single quotes. For example, use name = 'TCP' instead of name = TCP. (Fields should be named exactly as in the ArcSight Database when using wheresql. In this case, Name with uppercase would cause an error.)
- If **mathematical operators** such as square root or pi **contain a pipe**, the wheresql condition must be enclosed in double quotes. For example:

  |wheresql "bytesin > (|/ 25.0)"
- The "wheresql" condition field must exist in the current fieldset or be generated by a previous operator. For example, the field agentHostName must be contained in the *Base Events Field* fieldset. (Besides fields in a fieldset you can also use dynamic fields generated by previous operators. For example: | eval test1 = upper(name) | wheresql name != 'ArcSight')
- The "wheresql" condition cannot contain a limit. For example, the following statement is invalid: | wheresql Name = 'TCP' limit 1000
- Do not use the word "wheresql" for the name of a search, a search criteria, or a search query. The "wheresql" is a reserved word for the name of the search operator only.
- Do not use a semicolon at the end of the condition.

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

eval

The **eval** operator displays events after evaluating the results of a specified function. This can be a mathematical, string, or boolean operation and is evaluated when the query is run. The resulting value is assigned to a field name. Once a new field has been defined by the eval operator, it can be used in the query to further refine the search results.

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

General Syntax for Eval

The **eval** operator displays events after evaluating the result of the specified function. Eval operators use the following syntax formats:

```
| eval newField = expression
```

```
EXPRESSION Evaluation of values(fields) or constants with operators
```

where

- *expression* represents a valid field-based query expression.
- Arithmetic expressions or functions are not supported.

Functions that can be used with the eval operator include:

concat, tonumber, tostring, replace(X,Y,Z), abs(X), case(X,"Y",...), ceil(X), ceiling(X), exp(X), floor(X), if(X,Y,Z), isfalse(X), istrue(X), len(X), ln(X), log(X), lower(X), tolower(X), mod(x,y), rand(), round(X), sqrt(X), substr(X,Y,Z), sum(x,y,z,...), trim(X), ltrim(X), rtrim(X), upper(X)toupper(X), urldecode(X).

Aliases that contain special characters have the following syntax restrictions:

Special Characters	Restrictions	Examples
&, !, - , = , % , <, >, |	Need to be enclosed in single/double quotes when they are reused and the search works as expected.	| rename Name as 'DP-V' | eval test = tostring ("DP-V")
@, #, +, ?, /, ^, [], {}, _ , *, ., ~,\, $, .	Do not need to be enclosed in single/double quotes when they are reused and the search works as expected.	| rename Name as 'DP@V' | eval test = tostring (DP@V)
\	When a backslash is uses as part of an escape clause, add an additional backslash \. It does not need to be enclosed in single/double quotes when reused and the search runs as expected. The outcome field name should show only one backslash.	... | rename Name as 'DP\\V' | eval test = tostring (DP\\V)

For more information about eval functions, see "Understand Eval Functions" on page 111.

Considerations for Using Eval Functions

Please be aware of the following considerations when using the eval functions:

- You might encounter a search error if you run a query that uses both an "All Fields" fieldset and more than five pipeline operations. To avoid this, either reduce the number of fields in the fieldset or reduce the number of pipeline operators in your query.
- The md5(X) function is not supported in a FIPS environment.

Examples

- Could be a simple constant value: "Hello world", 5.

```
... | eval test0 = 5
```

- Could be a simple field: Name, Destination Hostname from the current selected fieldset.

```
... | eval test1 = Name
```

Pipeline operators, such as eval, can use operator chaining to allow output from one pipe operator to be used as input to a subsequent one.

- Find the longest URLs from the vendor ArcSight.

```
deviceVendor = ArcSight |eval urllength=length(requestUrl) |sort
urllength
```

- There is no limit to using arithmetic and boolean operators along with data (in fields or as constants).

```
... | eval test3 = Agent Severity + 1
```

```
... | eval test4 = Name and Device Vendor
```

```
... | eval test5 = (Name and Device Vendor) / 2
```

   If the boolean operator is the last operation applied, the overall result will be {0,1}.

- Examples of expressions that use arithmetic and boolean operators:

```
... | eval test6 = upper(Agent Severity ) + 1
```

```
... | eval test6 = upper(Agent Severity ) and Name
```

If the boolean operator is the last operation applied the overall result will be {0,1}.

- Example using "if" and "case" statements:

```
... | eval test = if ( deviceCustomNumber1 = 200, Success, Failure)
… | eval test = case ( deviceCustomNumber1 = 200, Success,
deviceCustomNumber1 = 400, Failure, Unknown)
```

- "Case" requires a final parameter that serves as the else condition. For example:

```
case(name = 'Mandy', 'analyst', name = 'Oskar', 'operator', unknown')
where 'unknown' is the else condition.
```

- Functions can receive other expressions as input:

```
...   | eval test6 = upper(Agent Severity and 1 ) and Name
```

Restrictions

Some functions have restrictions based on the data type:

The following expression is not allowed because two different data types (Name and 1) are not allowed in an arithmetic operation.

```
... | eval test1 = Name + 1
```

The expression below is not allowed because **replace** expects string data types for parameters.

```
... | eval test1 = Name and replace ( 1 , Name , Name )
```

For more information about syntax requirements that the query must meet, see "Understand the Query Syntax" on page 67.

Understand Eval Functions

Eval allows you to define and name an expression that is returned in the search. Use the following functions to build an eval expression:

For more information about other operators, functions, and syntax requirements, see .

Comparison and Conditional Functions

coalesce(*X*, [*Y, Z, N, ...*])

- Returns the value of the first non-null expression in the list. If all expressions evaluate to null, then coalesce returns null. The list is up to 20 elements long.
- In the list of expressions, all elements must be of same type.
- Parameters are the values used in the test.
- The only supported types are numeric and string. X can be a number, field or expression.

```
... | eval username = coalesce (Source Username, Destination Username)
Returns: Username
```

nullif(*X*,Y)

- Compares two expressions. If the expressions are not equal, the function returns the first expression (expression1). If the expressions are equal, the function returns null.
- *X* and *Y* can be a number, field or expression. *Y* must have same data type that *X*.

```
... | eval newField = nullif(2, 3)
Returns: 2
```

```
... | eval newField = nullif(2, 2)
Returns: null
```

Boolean Functions

and (&&), or (||), not (|), and **like**

- Results of boolean expressions are binary, meaning they satisfy conditions that are true/false, etc.
- You can connect and nest keywords. For example, (boolean_check_a) **and** (boolean_check_b).
- Use parentheses to group boolean operations.
- Each parenthesis should only do one binary "and/or" operation.
- Do not use more two boolean operators to connect keyword clauses. Instead, use parentheses to nest clauses. For example:

   Not allowed: (boolean_check_a) **and** (boolean_check_b) **and** (boolean_check_c)

   Allowed: ((boolean_check_a) **and** (boolean_check_b)) **and** (boolean_check_c)

```
| eval test_auto = (Agent Severity equals 4) and (Agent Severity equals
0)
```

```
| eval test_auto = (( Agent Severity equals 4 ) and ( Agent Severity
equals 0 )) and ( Agent Severity equals 2 )
```

Cryptographic Function

md5(*X*)

- Calculates the MD5 hash of string, returning the result as a string in hexadecimal.
- *X* must be a string.

```
... | eval usermd5 = md5 (Destination Username)
Returns: 202cb962ac59075b964b07152d234b70
```

> The md5(X) function is not supported in a FIPS environment.

Informational Function

isnull(*X*)

- Returns true if the *X* is null otherwise returns false.

```
... | eval newField = isnull(2)
Returns: false
```

Statistical Functions

greatest(*X,Y[,Z,N, ...]*)

- Returns the largest value in a list of expressions. The list is up to 20 elements long.
- In the list of expressions all elements must be of same type.
- The only supported types are numeric and string. *X* can be a number, field or expression.

```
... | eval newField = greatest(7, 5, 9)
Returns: 9
```

```
... | eval newField = greatest('sit', 'site', 'sight')
Returns: site
```

```
... | eval newField = greatest(bytesIn, 100)
Returns: 100, when bytesIn is less than 100
```

least(*X,Y[,Z,N, ...]*)

- Returns the smallest value in a list of expressions. The list is up to 20 elements long.
- In the list of expressions all elements must be of same type.
- The only supported types are numeric and string. *X* can be a number, field or expression.

```
.. | eval newField = least(bytesIn, bytesOut)
Returns: 5
```

```
... | eval newField = least('sit', 'site', 'sight')
Returns: sight
```

```
... | eval newField = least(bytesIn, 100)
Returns: 100, when bytesIn is greater than 100
```

randomint(*X*)

- Returns a random number between 0 and *X*-1.
- *X* can be any positive integer between the values 1 and 9,223,372,036,854,775,807.

```
... | eval newField = randomint(10)
Returns: a random number between 0 and 9
```

Text Functions

length(*X*)

- Returns the character length of a string, *X*.

```
... | eval n=length(field)
Returns: the length of (field). If the field is 256 characters long, it
returns n=256.
```

```
... | eval n=length("abc")
Returns: n=3 (abc is a literal string, surrounded by double quotes)
```

lower(*X*)

- Takes a string argument, *X*, and returns the lowercase version.

```
... | eval name=lower("USERNAME" )
... | eval name=tolower("USERNAME" )
Returns: the value of the field username in lowercase. If the username
field contains FRED BROWN, it returns name=fredbrown.
```

substr(*X,Y,Z*)

- This function returns a new string that is a substring of string *X*.
- The substring begins with the character at index *Y* and extends up to the character at index *Z*-1.
- The index is a number that indicates the location of the characters in string *X*, from left to right, starting with zero.
- *Y* can be negative.
- *Z* cannot be negative.

```
...| eval n=substr("ArcSight", 5, 6)
Returns: "g"
```

```
...| eval n=substr("ArcSight", 2, 6)
Returns: "cSig"
```

```
...| eval n=substr("ArcSight", 0, 3)
Returns: "Arc"
```

trim(*X*)

- trim(*X*) removes all spaces from both sides of the string *X*.

ltrim(*X*)

- ltrim(*X*) removes all spaces from the left side of the string *X*.

rtrim(*X*)

- rtrim(*X*) removes all spaces from the right side of the string *X*.

  For the sake of these following examples, assume that *X* is a literal string and _ represents any number of space characters.

  ```
  ... | eval trimmed=ltrim("_string_")
  Returns: trimmed="string_"
  ```

  ```
  ... | eval trimmed=rtrim("_string_")
  Returns: trimmed="_string"
  ```

  ```
  ... | eval trimmed=trim("_string_")
  Returns: "string"
  ```

upper(*X*)

- Takes one string argument and returns the uppercase version.

  ```
  ... | eval name=upper("username")
  ... | eval name=toupper("username")
  Returns: the value of the field username in uppercase. If username
  contains fred brown, it returns name=FRED BROWN.
  ```

For more information about syntax requirements that the query must meet, see .

concat

The **concat** function creates a new string field that concatenates (or links together) strings from fields. It concatenates any user-defined strings that are separated by a comma (",").

Search converts IP and MAC binary fields into a more user-friendly string format and then concatenated. Date fields are converted to the user format that is configured in your user preferences. Search converts NULL values to empty string fields. The maximum limit for concat results is 6,000 characters. Anything longer than this will be truncated.

- "Syntax" below
- "Parameters" below
- "How Do I Use This?" on the next page

Syntax

Syntax for concat should look like this:

```
... | eval newField = concat([field|value]*)
```

where:

- *newField* represents the field that you want to evaluate or test.
- *field* represents a string field in the result.

Parameters

The concat function can receive from 1 to 20 parameters, which can be expressions, user defined strings, or fields from the fieldset.

```
| eval test0 = concat('Event Name: ',  'Name')
```

```
| eval test1 = concat ( 'Event Name: ' , upper ( Name ) )
```

```
| eval test0 = concat ( 'Event Name: ' , ceil ( 2 ) )
```

```
| eval test0 = concat ( 'Event Name: ' , ceil ( 2 ) , replace ( Name , 'HTTP' , 'MQTT' ) )
```

How Do I Use This?

- Create an eval search that concatenates fields related to a Host:

```
| eval Host = concat(destinationHostName, ':' ,destinationPort)  - sample
output - mf.com:9000
```

- Create an eval search that concatenates fields related to the identity of an employee:

```
| eval Employee = concat(FirstName,' - ', LastName,  ' - ', DeptName,'(',
srcUserName, ')')
```

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

if and case

If() and **case()** are both eval operators that expect specified conditions be met (are True). An **if statement** returns a value when the given condition is met (is True), or returns another value when the given condition is not met (is False). A **case expression** runs through a set of given conditions and returns a value when the first condition is met (is True). Once the condition is met, the application stops any further searching for that condition. If no conditions are met (is False), then the software returns NULL.

When if() and case() are used together and the case expression is met (is True), the if statement returns that condition's values accordingly. If the case expression is not met (is False), then the application searches for the next condition given for the case expression.

- "Syntax" below
- "How Do I Use This?" below

Syntax

if

```
| eval test = if ( deviceCustomNumber1 = 200, Success, Failure)
```

case

```
| eval test = case ( deviceCustomNumber1 = 200, Success,
deviceCustomNumber1 = 400, Failure, Unknown)
```

The operators support conditional expressions, such as >, >=, <, <=, =, etc.

How Do I Use This?

Less than

if

- I want to determine if incoming bytes are less than 5000 and to identify the lowest and highest values of incoming bytes.

```
bytes in != null | eval test = if ( bytes in < 5000 , Low , High )
```

case

- I want to identify instances where incoming or outgoing bytes are below 3000 and return the lowest and highest values for each category.

```
bytes in != null AND bytes out is not null | eval test = case ( bytes in
< 3000 , Low , Bytes out < 3000, Low, High )
```

Equals

if

- I want to know all instances with an agent severity of 3.

```
agent severity != null | eval test = if ( agent severity = 3 , Success ,
Failure )
```

case

- Show me the device with the identification number 170011; otherwise show me the device with the identification number 3.

```
deviceCustomNumber1 is not null | eval test = case ( deviceCustomNumber1
= 170011 , Success , deviceCustomNumber1 = 3 , Failure , Unknown )
```

Contains numbers

if

- I want to identify all instances with a severity rating of zero (0) or one (1).

```
agent severity is not null  | eval test = if ( agent severity = 1 , 1 , 0
)
```

case

- Show me which devices have encountered a severity level of four (4); otherwise show me the highest and lowest severity levels.

```
agent severity is not null AND priority is not null | eval test = case (
agent severity = 4 , SHigh , priority > 5 , PHigh , other )
```

Three conditions

case

- I want to test three conditions (username, category outcome, and category technique) to identify Arcsight user names, any failed category outcomes, and any category techniques that might be exploited or represent vulnerabilities.

```
source username != null AND category outcome != null AND category
technique is not null | eval test = case ( source username = Arcsight ,
Arcsight , category outcome = '/Failure' , Failure , Category Technique =
'/Exploit/Vulnerability' , Vulnerability , other )
```

For information about other operators, functions, and syntax requirements, see "General Syntax for Eval" on page 108, "Understand Eval Functions" on page 111, and "Use an Operator in the Query" on page 85.

replace

The **replace** is a function of the eval operator that provides a mechanism to replace the content (expressed as string) of a [field](#) and to return the value in a new field. Before using replace, create a query that contains string values in its fields. When using replace, the process transforms the data into temporary tables so that the transformation occurs after the main query is executed.

- "Syntax" below
- "Parameters" below
- "How Do I Use This?" below

Syntax

```
Name is not null | eval test = replace(Name, "Response", "Returned Value")
```

where

- Name, Response, and Returned Value are the parameters used in the replacement function
- The replace function is case sensitive. For example, "This," "THIS," and "this" are considered three different words. Match the exact string in order to replace it

Parameters

Replace has three parameters:

- **Name**, the source string value
- **Response**, the match value that will be substituted with the returned value in the results
- **Returned Value**, the replacement value

How Do I Use This?

Use replace when you want to obfuscate data, improve the context of a field, or make reading the text more intuitive.

You can also use the replace function to replace an entire string.

- In this example, use replace to substitute a device's vendor name with Micro Focus.

```
| eval newDeviceVendor = replace ( deviceVendor, "HPE", "Micro Focus")
```

where:

- DeviceVendor is the source name for the string value.
- HPE is the response value.
- Micro Focus is the returned value.

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

tonumber

The **tonumber** eval function converts string fields into floating point numbers so that the data can be applied to additional calculations. If a result cannot be expressed as a number, Search leaves the field empty.

- "Syntax" below
- "Parameters" below
- "How Do I Use This?" below

Syntax

```
suboperator : tonumber
```

```
search_criteria | eval alias_name = tonumber one_field)
```

where:

- *search_criteria* represents a non-pipe operator query statement, such as "deviceVendor IS NOT NULL."
- *alias_name* represents a valid field alias.
- *one_field* represents a valid field as a parameter for "tonumber."

Parameters

There can be only one *fieldName* such as a device vendor or a version.

How Do I Use This?

Use tonumber to convert string values to numbers.

- Create a search query that converts log messages to numbers:

  ```
  | eval messagesAsNumber = tonumber ( message )
  ```

- Create a search query that converts vendor devices to numbers:

```
| eval x = tonumber ( deviceVendor )
```

- Create a search query that checks for vendor device data that is not NULL and convert the data from version fields to numbers:

```
deviceVendor IS NOT NULL | eval test = abs ( 10 ) + 10 | eval
toNumberAlias = tonumber(version) | eval test2 = abs (13)
```

- Filter the data for those entries where ArcSight is the device vendor. Transform the version to a number and the device's custom number to a string value:

```
Device Vendor = "ArcSight" | eval toNumberAlias = tonumber(version) |
eval numberToString = tostring (deviceCustomNumber1)
```

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

tostring

The **tostring** function is used in an eval operation to convert fields into string values. The input for tostring can be string values, numbers, integers, double point, float, IP/MAC address, and dates. All of these inputs must come from a field in the ArcSight Database.

- "Syntax" below
- "Parameters" below
- "How Do I Use This?" below

Syntax

```
search_criteria [pipe_operator]* eval alias_name = tostring (one_field)
[pipe_operator]*]*
```

where:

- *search_criteria* represents the criteria being tested in the query.
- *pipe_operator* represents the pipe operation for the query.
- *alias_name* represents the field to be converted to a string value.

Parameters

The function only accepts one parameter. More than that will cause an error. The parameter can be a field that represents a string, number, IP address, MAC address, and date. If the parameter is null, it returns a null input.

How Do I Use This?

Here are examples of queries using tostring:

```
... | eval testString = tostring(Name)
```

```
Name not equal null | eval testNumber = tostring(AgentSeverity)
```

```
... | eval testmac = tostring(Agent Mac Address)
```

```
... | eval testData = tostring(Device Receipt Time)
```

```
Agent Address not equal null | eval testIp = tostring(Agent Address)
```

For information about other operators, functions, and syntax requirements, see "Use an Operator in the Query" on page 85.

Understand the Search Criteria

You must have the **Manage Search Criteria** *permission.*

The **search criteria** defines the settings for your search, the time range in which to find data and the fieldsets that you want to use for displaying the results. Search provides "Understand the System Searches" on page 64 that you can view and load, such as DoS Events, MITRE ATT&CK Events, and Failed Login Event.

You can also save your search criteria for future use, such as loading the criteria into another search. You have the option to clone, modify, or remove a saved criteria at any time.

Manage the Fieldsets Displayed in Search Results

*You must have the **Create Fieldsets** permission.*

You can specify a **fieldset** that determines a group of search result fields the system displays in the Events table. In the table, each field can provide the ten most and less common values. Multiple searches can share a fieldset, and new searches display a default fieldset that contains the most common event fields. Use the fieldsets window to view and add the customize and system fieldsets, including lookup lists.

- **System Fieldsets**: Predefined fieldsets provided by the system.
- **My Fieldsets**: Customize the default fieldsets and lookup list fields for individual purposes.

New searches display the user´s default fieldset. These will remain selected in the fieldsets list box even when moving to other search tabs. If you select another fieldset, the pop-up window closes to display the new option. You can revert the change to the previously selected fieldset.

> Whenever you replace or update the fieldset, your search becomes out of sync, since the fields shown might differ from the new selection. Rerun the search with the new selection to correct this.

Create a Fieldset

1. From the **Search** page, click the icon to the left of the search name.

2. From the selected search's tab, click the menu and select a fieldset from the list in the **My Fieldsets**panel.

3. Click **Manage**.

4. Click **+** to add a new fieldset.

5. Enter a **Fieldset Name**.

   - Each fieldset should have a unique name.

   - Fieldset names are not case sensitive.

   - The fieldset is used only for your search results and does not affect other users connecting to the same system.

6. Select a **Category** and drag and drop any of the **Fields** to the to the **Selected Fields** column.

7. Click **Save**.

8. (Optional) Select **Apply to This Search** to customize the original fieldset without overwriting or saving it.

9. To execute the query again, click **Search**.

Edit a Fieldset

You can edit custom fieldsets only. You cannot modify system fieldsets, and you can only edit one fieldset at the time.

- "Editing the Selected Fieldset" below
- "Editing a Different Fieldset" below
- "Cloning a Fieldset" on the next page

Editing the Selected Fieldset

1. From the **Search** page, click the icon to the left of the fieldset name.

2. From the fieldsets window, click the fieldset menu and select a fieldset from the list in the My Fieldsets panel.

3. From the fieldsets window, select **Edit**.

   The **Edit Fieldset** window displays.

4. Drag and drop any field to the **Selected Fields** column OR select **Text Editor** to write the fields you need.

5. To locate a specific field, use the **Search** field.

6. In the **Fieldset Name** field, update the fieldset name as needed. The fieldset is used only for your search results and does not affect other users connecting to the same system.

7. Click **Save**.

8. (Optional) Select **Apply to This Search** to customize the existing fieldset without overwriting or saving it.

   The temporary fieldset will not be visible to other users, and it will only remain available on that session. After you log out, the system removes the temporary fieldset. You can have one temporary custom fieldset at a time.

Editing a Different Fieldset

1. From the **Search** page, click the icon to the left of the fieldset name.

2. From the fieldsets window, click the fieldset menu and select a fieldset from the list in the My Fieldsets panel.

3. Click **Manage**.

4. Select the fieldset checkbox.

5. Click the **Edit fieldset(s)** icon.

   The Edit Fieldset window displays.

6. Drag and drop any field to the **Selected Fields** column OR select **Text Editor** to write the fields you need.

7. To locate a specific field, use the **Search** field.

8. In the **Fieldset Name** field, update the fieldset name as needed. The fieldset is used only for your search results and does not affect other users connecting to the same system.

9. Click **Save**.

10. (Optional) Select **Apply to This Search** to customize the existing fieldset without overwriting or saving it.

    The temporary fieldset will not be visible to other users, and it will only remain available on that session. After you log out, the system removes the temporary fieldset. You can have one temporary custom fieldset at a time.

Cloning a Fieldset

You can make a copy of a fieldset you have created. Edit this copy to save you from creating a completely new fieldset.

1. From the **Search** page, click the icon to the left of the fieldset name.

2. From the fieldsets window, click the fieldset menu and select a fieldset from the list in the My Fieldsets panel.

3. Click **Manage**.

4. Select the fieldset checkbox.

5. Click the **Clone fieldset(s)** icon to make a copy of the selected fieldset.

Delete a Fieldset

You can delete a fieldset that you have created. If you delete a fieldset that's used in an active search, Search changes the fieldset name to **Custom** for that search. If you delete a fieldset used in a saved search query or saved search criteria, Search will use the default fieldset saved in your user preferences.You cannot delete a system fieldset.

1. From the **Search** page, click the icon to the left of the fieldset name.
2. From the fieldsets window, click the fieldset menu and select a fieldset from the list in the My Fieldsets panel.
3. Click **Manage**.
4. Select the fieldset checkbox.
5. Click the **Remove fieldset(s)** icon.
6. Click **Yes** to proceed.

Configure the Time Range

A search query can either have a fixed start and end date, where you cannot refresh data, or a time range that captures the most recent data. For example, if you choose the predefined **Last 30 minutes** setting, Search updates data upon re-executing the search based on the most recent 30 minutes. Alternatively, you can create a dynamic date range.

The time range that you specify in the time range selector is inclusive. Search includes the whole second as the end time. For example, if you specify a time range between *2018-01-01 12:00:00* and *2018-01-01 12:59:59*, Search includes all data from 2018-01-01 12:00:00.000 to 2018-01-01 12:59:59.999, inclusive.

- "Specify a Dynamic Date Range " on the next page
- "Understand the Search Timestamps for Events" on page 135
- "Understand How Time Zones Affect Search Results" on page 136

Specify a Dynamic Date Range

Search offers a flexible, dynamic setting for the time range where you can enter the desired time stamp without using the calendar to specify days, hours, and minutes. The dynamic date range uses the following syntax:

<dynamic_time>

or

<dynamic_time> [+/- *<units>*]

For example, to search for events that have occurred in the last two hours, you can specify `$Now - 2h` for **Start time** and `$Now` for **End time**. To find events that have occurred this week, you can enter `$CurrentWeek` for **Start time** and `$Now` for **End time**.

To enter a dynamic date range:

When viewing a search or starting a query, select the currently specified time range.

For the start or end time under **Custom Range**, select **Dynamic**.

To specify the **dynamic_time**, enter one of the following values:

Value	Represents
$Now	The current minute
$Today	Midnight of the current day
$CurrentWeek	Midnight of the previous Monday (or same as $Today if today is Monday)
$CurrentMonth	Midnight on the first day of the current month
$CurrentYear	Midnight on the first day of the current year

To specify the **units**, enter one of the following values:

Value	Represents
m (lowercase)	Minutes
h	Hours
d	Days
w	Weeks
M (uppercase)	Months

Understand the Search Timestamps for Events

Search can display results based on the timestamp associated with each event. The database stores three different timestamps for each event. For peak performance, Search automatically uses the Normalized Event Time setting. However, you can specify any timestamp setting for a search. You can also choose to make the timestamp the default setting.

> **NOTE**: The Date Picker displays this Timestamp setting when searching for events.

Database Receipt Time (dBRT)

Represents the time when the database received the event. The database considers this timestamp as the *persisted time* of the event.

Device Receipt Time (DRT)

Represents the time when the connected device claims the event occurred. This timestamp preserves the original time recorded by the device. However, this timestamp might not be credible in all cases. For example, it is possible that the time settings for the connected device are not configured correctly or the clock on the server that hosts the connected device might gain or lose time, which causes the timestamp to be out of sync with the actual time the event occurred.

Normalized Event Time (NET)

Represents the best known time for an event. Ideally NET is the time when the connected device reported that the event occurred (the DRT) because the device is the most direct known observer of the event occurrence. However, when the DRT for an event is not within a credible time range compared to the database's time, NET represents the time when the database received the event (the dBRT). For example, the time on a connected device was configured incorrectly such that DRT for an event is May 29 1975 when the current date in the database when the database received the event is June 29 2020. The database recognizes that the event's May 29 1975 timestamp for DRT is outside the credible time range. Based on the discrepancy with DRT, the database sets NET to June 29 2020 (same as the dBRT).

By default, the DRT value must be within a boundary of -7 days in the past and +1 days in the future from the dBRT. To configure the boundary criteria in a non-SaaS environment, see the *Administrator's Guide for the ArcSight Platform*.

Understand How Time Zones Affect Search Results

Searches for events in a time range are based on the timestamps of matching events and use the time zone of the local browser by default. You might need to account for the time zone offset from UTC and from other time zones, including Daylight Savings Time.

You can configure Search results to adjust the time for events to a specific time zone. For example, it's possible that you might create a search while in a one time zone, then view the search from a different computer set to a different time zone. When this occurs, the Events Histogram converts the time segments to the specified time zone. If the Events table includes a time attribute, Search converts the time. However, the aggregation reflects the original time zone. For example, if the Events Histogram has seven bars in the original time zone, the number of bars could increase or decrease to reflect the currently specified time zone.

Extend the Search with a Lookup List

Select **Configuration** > **Lookup Lists**.

You can make CSV files, or **lookup lists**, that enable Search to create additional tables with different fields and store them in the database. You can add lookup list fields to fieldsets and use them in search queries.

- "Understand the Considerations for the Lookup List File" on the next page
- "Create a Lookup List " on page 139
- "Append a Lookup List " on page 140
- "Replace a Lookup List " on page 141
- "Delete a Lookup List " on page 142

Understand the Considerations for the Lookup List File

The CSV file for your lookup list must meet the following requirements:

- The first row must be a comma-separated list of field names.
- The field names cannot exceed 40 characters. The names can only contain alphanumeric characters and underscores. They must start with an alpha character.
- For search operations, the corresponding data types for lookup lists with variable characters (or varchars) are short text (VarShort) and long text (VarLong).
- The remaining rows must be comma-separated values for the fields in the first row.
- Do not include spaces before, after, or within a field name.
- All rows must contain the same number of values.
- You must select one of the columns as the key field, and the values of the key field must be unique.
- The **key field** is the field that you can use with the `in list` operator in queries.
- The file cannot exceed 25 fields and 2 million rows.
- The file cannot exceed 150 MB.

Create a Lookup List

1. Select **Configuration** > **Lookup Lists**.

2. Click **Add**.

3. Drag-and-drop your CSV file to the **Lookup Lists** page or select **Browse** to navigate to the file.

4. Specify a name for the lookup list.

   Once created, you cannot change the name of the lookup list. The name must meet the following requirements:

   - Does not exceed 20 characters

   - Contains only alphanumeric characters and underscores

   - Starts with an alpha character

5. Specify the key field, then either accept the recommended value type or specify a different one.

   The following are possible values:

Value type	Specifies
domain	The name of the lookup list
float	A number whose radix point can be placed anywhere relative to the significant digits of the number
hostname	Fully qualified domain name
int	Integer value
ipv4	IPv4 address
ipv6	Ipv6 address
mac	MAC address
short text	Text that cannot exceed 1K of space
long text	Text that cannot exceed 4K of space
time	Time stamp
url	A URL address that cannot exceed 4K
username	A string type

6. To upload the file as a table in the database, click **Upload**.

Append a Lookup List

Use the **Append** feature to add more rows to a current lookup list.

- The file you need to append needs to have the same structure as the one you uploaded. For example, the same amount of columns.

- The file you need to append should not have an empty value in any of its rows.

1. Select **Configuration** > **Lookup Lists**.
2. Click the **eye** icon on the left side of the selected lookup list.
3. Click **Append**.
4. Select the list you want to append.
5. Click **Upload**. The original lookup list will be updated with the new rows added.

Replace a Lookup List

Replacing the contents of a lookup list does not affect queries that use the original lookup list. You cannot change the name of a lookup list. The field names in the replacement file must match the field names in the original file.

1. Select **Configuration** > **Lookup Lists**.
2. Select the list you want to replace.
3. Click the **eye** icon on the left side of the selected lookup list.
4. Click **Replace**.
5. Select the CSV file you want to use to replace the contents of the existing lookup list.

Delete a Lookup List

1. Select **Configuration** > **Lookup Lists**.

2. Select the list you want to delete.

3. Select the **trash can** icon.

Creating and Saving Searches

To execute a search, you must enter the query input, a fieldset that you want for the search results, and the time period for which you want to search events. Queries are case sensitive. The query input determines the search type (full text, natural language, or contextual). As you specify the search query, Search suggests search items and operators based on a schema data dictionary. You can also choose from predefined queries.

If you tend to use the same settings for some search parameters, you might want to configure your preferred default setting. For example, you can configure a default time range. To use the same search query or query plus criteria for multiple searches, you should **save** the query or criteria. You can also save the results of an executed search and configure a default expiration time for searches. By default, session searches expire after 24 hours of inactivity and saved searches after seven days. Search truncates long queries, displaying **…** to indicate additional content. To see the entire query, you can **pin** the input field.

> The application supports up to 10 active searches and 40 saved searches per user.

Create a Search

Select **Search** > **+**.

To execute a search, you must specify the query. You can use the default values for the fieldset, time range of data to search, and some additional settings or specify your preferred settings. Alternatively, you can load a saved query, criteria, or dataset.

If you tend to use the same settings for some search parameters, you might want to configure your preferred default setting. For example, you can configure a default time range. To use the same search query or query plus criteria for multiple searches, you should **save** the query or criteria. You can also save the results of an executed search and configure a default expiration time for searches. By default, session searches expire after 24 hours of inactivity and saved searches after seven days. Search truncates long queries, displaying **...** to indicate additional content. To see the entire query, you can **pin** the input field.

> If you exceed the search limit, the system displays following error message when you create a new search: "An error occurred while creating search. Exceeding the limit of 1000 searches." You cannot create anymore searches if this error displays. Contact your Administrator to increase the search limit or delete some existing searches. For more information about increasing the search limit in a non-SaaS environment, see Configuring the Deployed Capabilities in the *Administrator's Guide for the ArcSight Platform*. If you are a SaaS customer, reach out to Support to increase the search limit.

1. Select **Search** > **+**.
2. Enter the query in one of following ways:
   - To use a predefined System search, type **#**.

     The predefined searches might provide only a query expression or include search criteria such as a specific time range.
   - To use a search operator, such as eval and wheresql, begin typing the operator's syntax.

     For example, type:

     `... | where <expression>`
   - To manually enter the query, begin typing the expression.

     For example, type :

     `Source Address = 192.10.11.12 and Destination Address= 192.10.11.12 or Destination Address in Subnet 192.10.*.*`
   - To use a saved query, criteria, or search results, select 📁.

- To search data migrated from ArcSight Logger, select **Logger** from the list box next to the **Search** button.

- To search for a field without data, enter `[field_name] = Null`.

> In the query, Search treats a comma (,) between the search fields and values as an OR operator.

3. (Optional) To view all content in a very large query, select the **Pin** icon in the query input field.

   Otherwise, Search truncates long queries, displaying … to indicate additional content.

4. Specify the fieldset that you want for displaying the search results.

   By default, Search displays your preferred default fieldset. If you have not specified one, Search display the *Base Event Fields* fieldset.

5. For the time range, perform **one** of the following actions:

   - Accept the default time (**Last 30 minutes**).

   - From the menu, select a pre-defined value under **Quick Ranges**.

   - From the menu, use the **Custom Range** fields to specify a time range.

   - From the menu, select **Dynamic**, and then enter a dynamic date value.

   You can also specify the timestamp that you want to use for the retrieved events. Search uses "Normalized Event Time (NET)" on page 135 by default.

6. (Optional) To limit the number of results received from the search, complete the following steps:

   a. Select ⚙ to the right of the query input field.

   b. For **Maximum search results**, specify the maximum number of results that you want to receive in the dataset.

7. (Optional) If you do not want this search to expire in the default time, complete the following steps:

   a. Select ⚙ to the right of the query input field.

   b. For **Search expires in**, specify the number of hours that Search will store the session.

8. (Optional) To more easily find this session search later, give the search a name.

9. (Optional) To run the search, click **Search**.

   Alternatively, you can press **Enter** when editing the query input field.

10. (Optional) To save the query, criteria, or search results for future use, select the **Save** icon.

Load a Saved Search

If you have saved a query, criteria, or search results, you can load that saved item in a Search tab. You can also load the predefined search queries and criteria.

1. Select **Search** > **+**.
2. Select 📁 above the query input field.
3. Select the tab relevant to the saved search that you want to load:
   - Search Query
   - Search Criteria
   - Search Results
4. Select the saved search that you want to load.
5. Select **Load**.
6. (Optional) Modify the search settings as needed, then run the search.
7. (Optional) To more easily find this session search later, give the search a name.
8. (Optional) To save your changes as a new search, select 💾.

Run a Search

When you run a search, Search begins populating the Events histogram and Events table. Depending on the number of events retrieved, the search might pause to indicate that the amount of data could impact the search performance. You might want to select a smaller time range. To resume a search, click the play button in the progress bar.

1. Create or load the search that you want to run.
2. Click **Search**.
3. (Optional) To more easily find this search later, give the search a name.
4. (Optional) To save the query, criteria, or search results for future use, select the **Save** icon.
5. Click the **pause** or **stop** icons if you need to interrupt the search. Click the **resume** icon to continue the search.

Initiate a Search from Enterprise Security Manager

From Enterprise Security Manager (ESM), you can initiate a search in the ArcSight Platform for a maximum of five fields, based on the available columns on the active channel. Within ArcSight Platform, you can filter ESM data for more specific results. ESM generates a URL, opens a browser, and creates the new search in ArcSight Platform.

To perform this action, you must enable this feature in ESM. For more information, see the *ESM Installation Guide*.

Modify the Search Query or Criteria

When viewing a search, you can change the query, a fieldset, and the time range selection.

1. In the Search tab, change the query, fieldset, or time range.
2. To return to your original settings, select **Revert Changes**.
3. To update the search results with the modified settings, select **Search Now** or **Search**.

Name a Search

By default, Search gives each session search the title *Search <N>*. You can apply a custom name to the search at any time.

1. Right-click the name of the tab.
2. Select **Rename**.
3. Type the custom name.
4. Press **Enter**.
5. (Optional) To save the search, click the **Save** icon.

Search Event Data from Logger

Logger archived events can be viewed and consumed using the same parameters as in regular searches. From the **Search** page, hunt for ArcSight Logger events by selecting the Logger option from the list box next to the **Search** button.

> Before searching for Logger events from a particular Logger, metadata from that Logger must have already been imported, and at least one data migration from that Logger must have been completed, as described in "Importing Event Data From Logger" on page 328.

Before running a search on the Logger data, review the following considerations:

- Search supports only the specific set of operators available in the Search feature
- Your searches can include data from Logger storage groups even if the Logger storage groups do not display as part of the ArcSight Database's configuration.

1. Select **Search** > **+**.

2. From the list box next to the **Search** button, select **Logger**.

3. Add the required query details.
   You must use the search operators supported in ArcSight Platform.

4. Click **Search**.

> Note: If UTC time wasn't specified in the time range for importing events, you will need to convert the archive UTC timestamp shown in the **Import Logger Data** tab to your browser time/selected time zone, and enter that value as search time to fetch events from that time range

Save the Search

In any Search tab, select **Save** or in the saved list, click **+**.

You can save a search at any time. To save just the query or the query and criteria, you do not need to execute the search. After entering a query or criteria, or executing a search, complete the following steps:

1. In the Search tab, select the **Save** icon.

   Alternatively, when viewing a list of saved queries, criteria, or results, click **+**.

2. Select which part of the search you want to save:

   - Search Query

   - Search Criteria

   - Search Results (Dataset)

3. Specify a name for the saved search.

   - Each saved search must have a unique name.

   - We do not recommend using the same names for saved search queries, criteria, and results.

4. (Conditional) When saving the search results, specify how long you want to store the dataset.

   For example, if you have a Log Management and Compliance license and the *Never Expire Search Results* permission, you can configure a search to never expire.

5. Select **Save**.

Create, Update, and Navigate Searches Using a URL

Search lets you create, navigate, and update searches by modifying a URL. This is a quick way to integrate with other applications or to simply save your current search for future use.

Creating Searches

- Open Search and create a new search with default values.

  /rec/fusionSearch

- Create a new search with provided metadata.

  /rec/fusionSearch?query=Name+not+equals+null&fieldsetId=3223931b-b439-4951-a2d2-d83f9506c109&startDate=1673567387483&endDate=1673569187483

Opening Searches

- Open a specific search by the ID. The app will automatically add metadata to the URL after opening the search.

  /rec/fusionSearch/:searchId

  /rec/fusionSearch/home

Updating Searches

- Open a specific search by the ID and update the values. The app will ask you if your choice is the desired behavior.

  /rec/fusionSearch/:searchId?query=Name+not+equals+null

Opening a Search in the Event Inspector

- Open a specific event by GEID, ID, and Event Table in the Event Inspector tab.

  /rec/fusionSearch/eventsInspector?eventsDetail=
  [{"geid":":globalEventId","id"::Id,"eventsTable":":table"}]

> Sometimes, the timestamp for the DBRT (Database Receipt Time) is not selected correctly, as in the case for the URL *<FQDN>/re/search?query=Name%20<>%20null&timeColumn=persistedTime* This happens when you navigate to the URL from some (but not all) features in the system. Additionally, if you are not signed in and try to navigate to this URL, the search is not created, and the timestamp is not selected.

Learn more about how to create, save, run, and load saved searches by referring to their respective topics.

Viewing and Managing Your Searches

Search displays results in an **Events Histogram**, **Events** table, and **Event Details panel**. If connectors are configured to send raw events, the table and details panel can include **raw event data**. Also, the maximum number of events that a search can return is 10 million. If your searches regularly stop at the maximum limit, consider splitting the query into separate searches.

Get an Overview of Your Searches

Select **Search** > **Home**.

The **Search Home** tab provides a high-level view of your Search and event activity while also offering immediate access to search features. Your Search and event activity are segmented into Widgets that show the state of specific searches and events, such as saved search queries, search criteria, and events by agent type. To view the items referenced in a search chart, such as saved search queries, click ⬀.

> The data in the events pie charts (Total Events, Total Events - 24 Hours, Events by Device Vendor - 24 Hours, etc.) refreshes every five minutes and is based on a normalized event time.

My Session Searches

Lists all recent searches. The Session Searches table includes the following columns for each search: Name (sortable), Search Query, Search Status (sortable), Fieldset, Start Time, End Time, Executed (sortable). To view one of the searches, click 🔍. To remove session searches from the list, select the rows to be removed. Then, click 🗑.

> The Session Searches table will only show searches with the status Completed, Pause, Running, and Error.

Search Queries

Shows the number of system, private, and public saved search queries that you can access.

Search Criteria

Shows the number of system, private, and public saved search criteria that you can access.

Search Results

Indicates whether any of your saved search results have completed, are running, or have been paused.

Fieldsets

Shows the number of system and private fieldsets that you can use when running a search.

Lists

Shows the number of lookup lists that you can include in a search.

Total Events

Shows the total number of all events that have occurred within your ArcSight database installation. This number refreshes every five minutes.

Events 24 Hours

Shows the total number of all events within the last 24 hours.

Events by Device Vendor - 24 Hours

Shows the total number of the first three events by device vendor with the highest count within the last 24 hours.

Events by Agent Type - 24 Hours

Shows the total number of the first three events by agent type with the highest count within the last 24 hours.

Events by Agent Severity – 24 hours

Shows the total number of the first three events by agent severity with the highest count within the last 24 hours.

View the Results of a Search

Search results are displayed in an Events Histogram, Search Results table, and Event Inspector panel. If connectors are configured to send raw events, the table and inspector panel can include raw event data. Also, the maximum number of events that a search can return is 10 million, but you can specify a preferred limit. If your searches regularly stop at the maximum limit, consider splitting the query into separate searches.

You can export the search results to a CSV file.

- View the Event Histogram
- View the Search Results Table
- View the Event Inspector

View the Event Histogram

The Histogram displays data in a segmented graph where the y-axis presents the number of events per bars of time segments in the x-axis. The time range on the x-axis might not match the time range specified in the search query because the start and end times on the x-axis are determined by the event times of the first and last matching events of the search query.

Click the menu to the right of the histogram and select either Linear Scale or Log Scale to display the data in your preferred format. As you hover your pointer over the histogram, the bar color directly below the pointer changes and displays a tooltip of the day/date/time of that event range. Click a bar to view event information for a specific time range. Click again to deselect the bar.

Note that some search activities do not require the histogram, and thus it will not be displayed. For example, if you perform an aggregation operation, such as "top" or "bottom," Search will not display the histogram because the Search Results table contains the aggregation of results, not events in a timeline.

How Search builds the histogram

Search progressively builds the histogram as it receives events that match the search settings. If the search needs to scan a large amount of data or a large time period, the histogram displayed initially might refresh multiple times while the search is running. To view the complete histogram of a search, wait until the search has finished running.

Search plots the first one million matching events on the histogram. If a search results exceed one million events, Search displays an informational message. If you need to use the histogram view for event analysis of a search that matches more than one million events, we suggest that you adjust the time range to retrieve fewer than one million events. This will allow you to obtain a complete and meaningful histogram. You can also use a pipeline operator to further refine search results so that the total number of hits is under one million events.

Narrow the scope of the search

If you have a large number of data points or a wide time range, you can see the big, overall picture, but you might not be able to clearly identify specific data points. To narrow the scope of the displayed data, adjust the boundaries of the displayed bars. As you adjust the time range within the Histogram, the Events table displays corresponding events.

Drill down to events

You can drill down to events in a specific time period by clicking the bar on the histogram that represents that time period. The bar you drilled down to is highlighted and the events matching that time period are listed below the histogram. To deselect the time period, click the bar again. When you **hover over a histogram bar**, the matching events listed below the histogram do not change, and the histogram continues to display all matching events.

View the Search Results Table

The **Search Results** table contains all the fields specified in the fieldset. You can choose to display the table in **Grid View** or **Raw View**. You can perform the following actions while viewing the table:

View all details for an event

To view details of a specific event, right-click the event and select **Open In Event Inspector**. This action opens the Event Inspector in a panel on the right where you can view additional details on the event.

View raw event data

When you click the **Raw View** icon, the Search Results table replaces the fieldset with a Raw Data column, which displays the whole raw event. Although the **Raw Event** field is most applicable for syslog events, you can also display the raw event associated with CEF events.

To do so, make sure the connector that is sending events to the database populates the *rawEvent* field with the raw event.

Export the search results

To export the results to a CSV file, select ⬇.

Export a single event

To export a single event, right-click the event. Then, select either **Export to PDF** or **Export to CSV**.

Copy a value from an event

To use a value from an event elsewhere, simply right-click and copy the value.

Compare data in columns

Hover over a column heading, then click the **Pin** icon to pin or unpin a column.

By pinning a column, you can compare the column's values against those of other columns. Search moves the pinned column to the extreme left location in the table. You can pin multiple columns.

Reorder columns

To rearrange the order of the columns, drag each column to new position by clicking and dragging the column header.

Sort the data in columns

Select the **up or down arrow** in the column heading to change the sort order.

View the Event Inspector

The Event Inspector displays additional details on any event you select from the Event table. This panel allows you to scroll through the specific details of the event and groups the details by categories such as **Agent** and **Source**. To open the Event Inspector, right-click any event in the Search Results table. Then, select **Open in Event Inspector** from the pop-up menu.

> To view events migrated from Logger, select **Logger** before creating a search.

You can perform the following functions with the Event Inspector:

Search for fields and values

To search for fields and values in the details of an event, enter a string in the search box at the top of the Event Inspector. The Event Inspector will filter the fields and values to match your search criteria.

Add fields and values to current or new search

You can add event fields and values to your current search or a new search.

Hover over a field (for example, Agent Hostname) to display a check box next to the field. Then, select the check box to select the field and its value. Then, either click the magnifying glass icon at the top of the Event Inspector or right-click your selected field. Both actions display a pop-up menu with the following options:

- **Create New Search:** Selecting this option allows you to create a new search query with the selected event fields and their values. For example, if you selected the field "Name" and its value equals "failed login", then it would display as follows in the new search query: Name = failed login. The new search will open in a new tab on your web browser.

- **Add to Active Search:** Selecting this option adds your selected event fields and their values to the current search query in the search input field. For example, if you selected the field "Name" and its value equals "failed login", the field and value would display as follows in the current search query: <current search query> | where Name = failed login.

Copy and share event detail URL

To share event details with another Analyst, click the **Copy URL** icon at the top of the Event Inspector. This action copies the URL to your clipboard so you can share it as needed.

Export event details to PDF or CSV

To export event details to a PDF or CSV format, click the **Export** icon at the top of the Event Inspector. A pop-up menu opens with the options **Export to PDF** and **Export to CSV**. Select the option that best meets your needs. You can include or exclude null fields in the exported file.

Expand/collapse and show/hide data fields

The top of the Event Inspector contains an arrow icon that expands and collapses the event details. There is also an eye icon that can show or hide null fields. If you select to display null fields and export the event details to PDF or CSV, the exported file will contain the null fields.

View and Use the Details of an Event

Right-click an event in the Search Results Table > click **Open In Event Inspector**.

The Event Inspector opens in a panel that allows you to scroll through the details of an event and groups them by categories such as **Agent** and **Source**. Use this panel when you want to research specific details on an event.

You can view the raw data details for the event, as well as instruct the panel to include fields with *null* data. For example, you could view details about the agent, category, device, source, or severity. You can only open one event in the Event Inspector at a time.

> To view events migrated from Logger, select **Logger** before creating a search.

- Search for Event Details
- Copy and Share Event Detail URL
- Export Event Details to PDF or CSV
- Apply Event Details to Current or New Search
- View Null Data Fields
- Expand or Collapse All Data Fields

Search for Event Details

The top of the Event Inspector contains a search box that allows you to search through the fields in the event details. Use this feature to quickly locate specific details on an event without the need to scroll through the entire Event Inspector.

To search for fields and values in the details of an event, enter a string in the search box at the top of the Event Inspector. The Event Inspector will filter the fields and values to match your search criteria. For example, if you searched the term "device" the panel will display all fields with the name "device" and any fields containing the value "device".

Copy and Share Event Detail URL

You might want to share the selected event's details with an Analyst or use the details in a report or other media. You can export all content in the Event Inspector with or without empty values.

Click the **Copy URL** icon at the top of the Event Inspector to copy the Event Inspector URL to your clipboard. Then, you can share the URL as needed. When an Analyst loads the URL, the Event Inspector will open in their browser with the event details related to the URL.

This action is helpful in situations where you need an Analyst to research an event further or for reporting purposes.

Note: The Event Inspector URL contains the event's ID (id field in the Search Results table) and global event ID (geid field in the Search Results table). See the table below for an example and variations of the Event Inspector URL format. Use these formats to create the URL.

> If the geid is missing in the URL, an error message will display.

Event Inspector URL	Example
Full Event Inspector URL	/rec/fusionSearch/eventsInspector/?eventsTable=Recon&id=5139791690&geid=3009625190352082178
geid and id only	/rec/fusionSearch/eventsInspector/?id=5139791690&geid=3009625190352082178
geid only	/rec/fusionSearch/eventsInspector/?geid=3009625190352082178

Export Event Details to PDF or CSV

There may be situations where you need to use event details for reporting purposes. Or, you may need to share the event details with an Analyst who does not have access to the Event Inspector. You can do so by exporting the event details to PDF or CSV. Follow these steps:

1. At the top of the Event Inspector, click the **Export** icon.

2. A pop-up menu appears. Click either **Export to PDF** or **Export to CSV**.

3. Both selections will start a download of the event details to your selected format.

4. Share or use the PDF or CSV as needed.

If the option to show null values is selected, those null values are included in the exported CSV or PDF file. If null values are excluded, they will not appear in the exported file.

NOTE: You can also export an event to PDF or CSV from the **Search Results Table**. Right-click an event in the Search Results table to open a pop-up menu with the options **Export to PDF** and **Export to CSV**. If you use this method to export the event details, null values will be included in the exported file.

Apply Event Details to Current or New Search

You can add the field and value pairs in the event details to your current search or a new search. This action is helpful in situations where you need to research more data on a specific event. After adding a field and value pair to a current search or new search, you might need to add the respective field to the search fieldset if that field is not already part of the fieldset.

Hover over a field in the Event Inspector (for example, Agent Hostname) to display a check box next to the field. Then, select the check box to select the field and its value. From here, do one of the following actions:

- Right-click the selected event field
- Click the magnifying glass icon at the top of the Event Inspector

Both actions display a pop-up menu with the following options:

- **Create New Search:** Selecting this option allows you to create a new search query with your selected event fields and their values. For example, if you selected the field "Name" and its value equals "failed login", then it would display as follows in the new search query: | where Name = failed login.

- **Add to Active Search:** Selecting this option adds your selected event fields and their values to the current search query in the search input field. For example, if you selected the field "Name" and its value equals "failed login", the field and value would display as follows in the current search query: <current search query> | where Name = failed login.

Once you've performed a new search with the selected field and value pairs, the Event Timeline and Search Results table will filter to display data related to your new search.

View or Hide Null Data Fields

To show or hide fields with null data, click the eye icon at the top of the Event Inspector. Hiding the null fields filters your view of the event details to show only fields with data. Use this feature if you want to see only fields with data in the event details.

Expand or Collapse All Data Fields

Next to the eye icon at the top of the Event Inspector is an **Expand All/Collapse All** icon. Click this icon to expand the fields in the Event Inspector to show all values related to the fields. Or click it to hide the values related to the fields and display only the field names.

Identify Fields without Data

If an event does not have data for a schema field, Search represents the absence of data (*null*) in the results in the following ways:

Affected Field	Displayed Result
Search field	Null, NULL and null query formats
Events table	Empty cell
Empty field from ESM (for example, name='')	name = '', NULL
Event Details panel	--- in the cell

Refresh Search Results

If the time range for your search is based on a predefined range, such as **Last 30 minutes**, you can refresh the search results as desired. However, refreshing the browser as you update a search does not save your changes. You must save the refreshed results.

Export the Search Results

You can export the Events table to a CSV file. Search exports data based on the specified fieldset for the search. The export process limits the file to one million event records.

1. In the table's header, select the **CSV** icon.

2. In the search toolbar above the histogram, select ⬇.

3. Choose to save the file or open in a desired application.

> Saved .csv files of search queries sometimes contain syntax that uses "name= a special character" (such as "name=+" or "name=_"). In Excel, if the file is not opened properly, you might see formatting issues where the fields display as a generic "#NAME".
>
> To avoid this, open the file by selecting **From Text/CSV** on the Excel **Data** ribbon. Navigate to the .csv file you downloaded and click **Import**. Preview the file to ensure the fields display correctly, then click **Load** to view the full file.

Configure Preferred Settings for Searches

Select **[your_ID]** > **My Profile** > **Preferences**.

You can specify the default settings that you want to apply for new searches. For example, you might want all of your searches to return results from the last 24 hours. Or, if you regularly use the same fieldset for a Search, you can specify that fieldset as your preferred default. You can always override your preferences as needed when you create a search. When you modify your Search preferences, the changes apply to new searches. Existing searches are not affected unless you re-run the search.

> If you change your search preferences and you also have Scheduled Searches open in a separate browser tab, you must refresh the Scheduled Searches tab to ensure that the content in the tab reflects your changes.

Default Fieldset

Specifies the fieldset that you regularly use for a search. The default value is *Base Event Fields*.

Default View

Specifies whether the Events table displays results in the Grid View or Raw View. The default value is *Grid View*.

Time Zone

Instructs Search to adjust the timestamp for events to the chosen time zone.

Date/Time Format

Specifies the format of dates and times you want Search to use. The default is *MM/DD/YY hh:mm:ss:ms*.

Default Time Setting

Specifies the time range you want Search to find events. The default is the *Last 30 minutes* Preset value.

Base Searches On

Specifies the timestamp Search associates with the event you want to find. The default value is *Normalized Event Time*.

Search expires in

Specifies how often you want saved searches to expire, and thus for the system to remove them from the system. You can specify a value between 1 and 365. The default value is 7 days. Alternatively, if you have the *Never Expire Search Results* permission, you can choose for a search to never expire. When you create or edit a search, you can override this default setting.

The expiration date resets whenever you access the search. Resetting the expiration date includes resuming or re-running the search, as well as saving the search and changing its settings.

Session Search expires in

Specifies how often you want session searches to expire. The default is *24 hours*. You can specify up to 120 hours. The expiration time resets whenever you change or run the search. When you create or edit a search, you can override this default setting.

Maximum search results

Specifies the maximum number of events that Search returns. Search considers a search complete when the results reach the maximum limit. The default value is *10,000,000*. The lowest value that you can specify is 1,000. When you create a search, you can choose to override this default setting.

Your admin can configure a system-level setting that controls the maximum number of searches (with a limit of 10 million) for all instances of Fusion. If you enter a value outside of the system-level setting, you will receive an error message indicating that your preferred default cannot exceed the system setting. For information about setting a global search limit, see Upgrading Deployed Capabilities in the *Administrator's Guide to ArcSight Platform*.

Highlight Query Syntax

Specifies whether Search uses color to differentiate the syntax terms from the operators and functions within the query. The default value is set as *Yes*.

Manage Searches

Select **Search** > **Search** *<saved_search_type>*.

If you have saved a search query, criteria, or dataset, you can manage the saved items individually or in bulk, import and export search results in a CSV file, or delete them.

Manage Your Search Queries

Select **Search** > **Search Query**.

The saved search queries contain only the specified query expression, ready for you to load into a new search at any time. The list of saved queries includes both the queries that you have saved and the built-in System queries. You can **modify** or **delete** your queries at any time. However, you cannot delete or edit a System query. Rather, to change a System query, you should clone it, then make and save your changes.

Import a search query (as a gzipped JSON file) by clicking the `Import` icon, selecting the desired file, and clicking `Import`.

An imported file cannot exceed 100 MB, must contain only search queries with valid information.

You can also **export** one or more queries to a gzipped JSON file.

You must have *Import and Export Search Queries* permission to either import or export queries.

Manage Your Search Criteria

Select **Search** > **Search Criteria**.

Saved search criteria combine a query expression and other Search elements such as fieldsets and the time range of the data you want to retrieve. The list of saved criteria includes both the criteria that you have saved and the built-in System criteria. You can **modify** or **delete** your criteria at any time. However, you cannot delete or edit a System criteria. Rather, to change a System criteria, you should clone it, then make and save your changes.

By default, search criteria are sorted alphabetically by name. Date fields are displayed according to your user preferences.

If you have the *Import and Export Search Criteria* permission, you can **import** or **export** one or more criteria to a JSON file.

Import

1. Select **Search** > **Search Criteria**.
2. Click the **Import** icon.
3. Select the gzipped JSON file (or files) you want to import.
4. Click the **Import** icon.

   The selected criteria (and any associated fieldsets not already in the system) are imported.

Export

1. Select **Search** > **Search Criteria**.
2. Select the entries that you want to export.
3. Click the **Export** icon.

   The selected entries download into a gzipped JSON file.

Manage Your Search Results

Select **Search** > **Search Results**.

When you save search results, Search stores the dataset until the search expires or you delete it from the saved list. You can sort the list by the search's name, query, event time stamp, or date of the search.

If you load the saved dataset in a Search tab, you can update the query and criteria as needed, then save those changes as a new search query, criteria, or results. To share the results with colleagues, export the results to a CSV file.

Scheduling Regular Runs of a Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Schedule**.

A **scheduled search** is a search that runs on a regular interval. Whereas a saved search is saved, but does not run automatically. Each time a scheduled search runs, search adds the results to the list of Completed Searches runs.

Manage Scheduled Searches

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Scheduled**.

- "Create a Scheduled Search" on the next page
- "Manage Scheduled Searches" above
- "Clone a Scheduled Search" on page 178
- "Edit a Scheduled Search" on page 179
- "Delete a Scheduled Search" on page 181
- "Enable and Disable a Scheduled Search" on page 180

For your scheduled searches, you can perform the following actions:

View and edit all details for a schedule search
To view specific scheduled search details, in the Name column, locate the search name and select it. Click **Edit** at the top of the table.

Sort the data in columns
To change the sort order, click the column heading to toggle between ascending and descending order.

Reorder columns
To rearrange the order of the columns, drag each column header to a new position.

Search for a search keyword
To find a keyword, click the field next to the **Magnifying Glass** icon (Search Keyword), enter a value, and the system displays your results automatically.

Hide and display columns
To hide and display a column, in the far right-corner of the window, click the **Wrench** icon (Manage Columns), and then select and clear the column name checkboxes.

Filter the data in columns
You can filter scheduled searches based on Status, Timestamp, and Fieldset. To filter the data for more specific results, in the far-right corner of the window, click the **Funnel** icon (Filters), and then select and clear the filter options.

Create a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

For every scheduled search, enter the query, fieldset, or time range for the search events or leave the defined values for the saved search. Just as for a saved search, the following considerations apply to a scheduled search:

- The search is case sensitive.
- The query input determines the search type (full text, natural language, or contextual).
- The system treats a comma (,) between search items and values as an OR operator.
- As you specify the search criteria, the system suggests search items and operators based on a schema data dictionary. To view the predefined queries, type **#** in the **query** field.
- To search for a field without data, enter [field_name] = *Null*.

To create a scheduled search:

1. Select **Search** > **Scheduled Searches**.
2. Select **+**.
3. Specify a **Name** that is 5 to 255 character long.
4. To enable the scheduled search, select **enable**.

   You also can enable and disable scheduled searches at any time in the **Scheduled** tab.
5. To indicate how frequently you want the search to run, specify one of the following options:
   - **Hourly**
   - **Daily**
   - **Weekly**
   - **Monthly**
6. Configure the settings for the dates and times of each run, based on how frequently they will run.

   > **NOTE:** If you choose the **End after** option, the maximum number of instances is 1000.

7. For **Search Query and Metadata**, complete one of the following actions:
   - To use an existing search, type # then select from the list of available saved searches.

- To create a new search, specify the query, fieldset, and time range.

8. Select **Schedule**.

Clone a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Scheduled**.

After creating a scheduled search, you can clone it at any time.

1. Select the scheduled searches that you want to clone.
2. Click the **clone** icon.

Edit a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

After creating a scheduled search, you can edit it at any time. After you modify a schedule, the first completed run will have a flag to indicate that the modification occurred.

If you change the **Pattern** values, please be aware that Search counts any and all completed runs before you made the change. For example, your scheduled search uses the **repeat forever** option and Search has performed three runs. If you update the **ending option** to end after eight occurrences, Search counts the three previous completed runs; therefore, you would only have five occurrences of the eight occurrences left to run. Should you want eight occurrences, you would need to change your **ending option** to 11 occurrences.

1. Select **Search** > **Scheduled Searches**.
2. Select the scheduled searches that you want to edit.
3. Click the **edit** icon.

Enable and Disable a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Scheduled**.

After creating a scheduled search, you can enable and disable it at any time.

1. Select the searches that you want to enable or disable.
2. Select **Enable** or **Disable**.

    The **Status** column, which you can add with the *Manage Columns* option, displays the status of either **Enabled** (green) or **Disabled** (red).

Delete a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Scheduled**.

You can delete a scheduled search at any time. After selecting **Delete**, the system prompts you to keep or delete the completed runs associated with the scheduled search.

> To cancel the deletion process, select the **X** that closes the dialog box, instead of selecting **Yes** or **No**.

Manage Completed Runs of a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Completed**.

After creating a scheduled search, you can view, delete, export, and filter the **completed runs** of that search. The results of a completed run are immutable. That is, if you edit the settings or query of a completed run, your changes do not affect the original results stored in the Completed list of scheduled searches.

- "View a Completed Run of a Scheduled Search" on the next page
- "Save the Results of a Completed Run" on page 185
- "Delete Completed Runs of a Scheduled Search" on page 186
- "Export Completed Runs of a Scheduled Search" on page 187

View a Completed Run of a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Completed**.

The name of a completed run represents the name of the scheduled search name plus its start date and time.

When a run is in progress, Search displays the number of events received thus far and when the last chunk of data was received. Also, a flag beside the name of a completed run indicates that the settings for that scheduled search were changed before this run.

> A completed run can fail if the global search limit is exceeded. To verify if the global search limit was reached, create a new search. If the error message: "An error occurred while creating search. Exceeding the limit of 1000 global searches." displays, then the limit was reached. Contact your Administrator to increase the search limit or delete some existing searches. If you are a SaaS customer, reach out to Support to increase the search limit. For more information about increasing search limits, see Configuring the Deployed Capabilities in the *Administrator's Guide to ArcSight Platform.*

In the **Completed** tab, you can perform the following actions:

View all details for a completed schedule search

To view completed search results, click the **Eye** icon beside the search name.

Sort the data in columns

To change the sort order, click the **column heading**.

Reorder columns

To rearrange the order of the columns, drag each column to new position.

Search for a search keyword

To find a keyword, click in the field next to the **Magnifying Glass** icon (Search Keyword), enter a value, and the system displays your results automatically.

Hide and display columns

To hide and display a column, in the far right-corner of the window, click the **Wrench** icon (Manage Columns), and then select and clear the column name checkboxes.

Filter the data in columns

To filter scheduled searches based on *Status* and *Fieldset*, select the corresponding filter parameter. You can also filter completed scheduled searches based on a time range (custom and preset).

To filter the data for more specific results, in the far right-corner of the window, click the **Funnel** icon (Filters), and then select and clear the filter options. To filter the results based on execution time, set the date picker filter in the far right corner.

Save the Results of a Completed Run

Select **Search** > **Scheduled Searches** > **Completed**.

You can save the dataset from the completed run of a scheduled search, similar to saving other searches. When you save the run results, Search renames the selected run to the name that you specify. You also can choose how long to retain the dataset in the database.

1. When viewing a completed run, select the **Save** icon.
2. Specify a name for the saved dataset.
3. Under **Result Retention and Limitations**, configure how long you want to keep each completed run of the scheduled search.
   - Your choice of values for each setting might be confined to limits set by your product administrator.
   - For **Delete files after**, you can specify a value that overrides how you configured Search Expires In for your search preferences.
     For example, you prefer that searches expire within five days. But you want the dataset for this completed run to expire after 10 days.
   - (Conditional) If you have the *Never Expire Search Results* permission, you can choose **Never Expire** to retain the dataset indefinitely.
4. Select **Save**.

Upgrading to the New Search Capability

After you upgrade to the new Search capability, you might encounter minor issues with saved scheduled searches. The general workaround to prevent these issues is to save your previous results **before** the upgrade and recreate them for new search runs. Issues you might see include:

- For completed scheduled searches before and after an upgrade, the "number of results" column may not match actual search results or equal zero. But, you can still view the actual results by opening the completed scheduled searches.
- The results of scheduled searches that contain the **eval** operator may not load properly if you are loading them in a search results tab that is already open.

Delete Completed Runs of a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Completed**.

You can delete a completed run of a scheduled search at any time.

1. Select the completed runs that you want to delete.
2. Click the **delete** icon.

Export Completed Runs of a Scheduled Search

*You must have the **Manage Scheduled Searches** permission to schedule runs of a search.*

Select **Search** > **Scheduled Searches** > **Completed**.

You can export the completed run of a scheduled search to CSV format.

1. Click the **CSV** icon next to the name of the scheduled search that you want to export.

2. Alternatively, view the search, then select the **CSV** icon to export the results.

Checking the Integrity of Event Data

*You must have the **Perform Event Integrity Check** permission to run a check.*

Select **Admin** > **Event Integrity**.

The ArcSight Database stores all collected events to support event searches and analysis capabilities for the ArcSight Platform. When investigating a security incident or hunting for threats, you expect that the search results provide valid and accurate data. However, the data that you rely on could be compromised by individuals who want to hide their activities or who maliciously change content. Data also is vulnerable to human errors, transfer errors, or loss and corruption caused by hardware or software issues. To reduce the chance of data tampering, the database enforces the immutability of events once they are stored, ensuring that not even the most privileged database administrator can modify or delete an event. You can also run an **Event Integrity Check** to validate that the event information in your database matches the content sent from the SmartConnectors. The combination of an integrity check with the database's ability to resist tampering provides you an end-to-end, long-term solution for safeguarding the events to be exactly as reported by the device where the activity was observed.

When you run the check, the system searches the database for verification events received within the specified date range, then runs a series of checks to compare content in the database with information supplied by the verification events. The results of an Event Integrity Check help you identify whether event data might be compromised or incomplete. The event integrity checks can involve two different types of verification events: generated for *raw events* from SmartConnectors or for *parsed fields* from Transformation Hub. Both types of verification events can be used in the same environment for increased visibility into the integrity of the events in the database.

> **NOTE** : At this time, the Event Integrity Check searches only the events ingested from SmartConnects to the ArcSight Database and does not include events migrated from Logger.

Understand the Event Integrity Check

Depending on how you have configured your SmartConnectors and Transformation Hub, the Event Integrity Check can verify raw event data and the parsed fields within an event, respectively. The check looks for events referenced by verification events in the database. The SmartConnectors group several events and compute a hash for each raw event in the batch. If you use Transformation Hub as a destination, it also groups events then generates a hash for the parsed fields within each event. The SmartConnector and Transformation Hub each generates a *hash of the individual hashes* to create a **verification event**. The number of events in a batch depends on how you configure the batch size setting for each connector. Note that SmartConnectors do not store the hashes for individual events.

Figure 1 (below) shows how events flow from your data sources to the SmartConnectors, which generate the verification events for the raw events. Then Transformation Hub generates verification events for parsed fields within each event.

Figure 1. *Process for generating verification events for an Event Integrity Check*
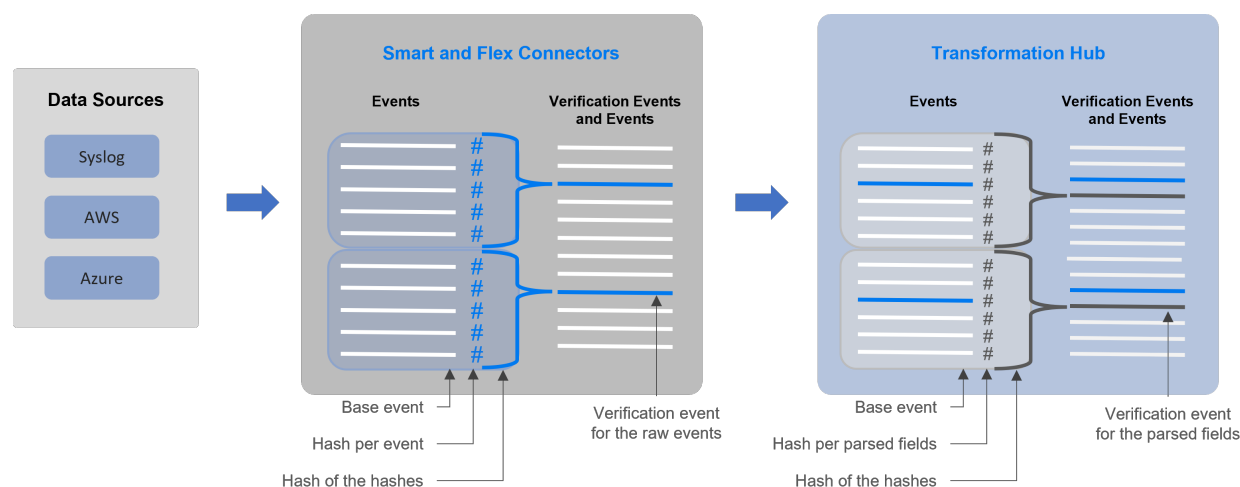
Each verification event includes the following items:

- a group of events with raw data or events with parsed fields
- an ordered list of the event IDs within the batch
- a crypto signature field representing the computed hash for that batch (the hash of hashes)

When you **run an Event Integrity Check**, the system performs the following actions for each verification event in the specified time range:

- Looks for the globally unique event ID (GEID) of each event referenced with the verification event.

- Generates hashes for the events within the base event.

- Generates a hash to represent the base events' hashes in the sequence provided by the verification event. You might call this the *generated hash of hashes*.

- Compares the generated hash of hashes to the hash of hashes in the crypto signature field that the SmartConnector or Transformation Hub created for the verification event.

> Some base events could have been deleted on purpose to comply with data retention policies, depending on storage group configuration. When performing an event integrity check, the system reports these deleted events as missing base events.

Run an Event Integrity Check

An Event Integrity Check looks for verification events received within the specified date range. To reduce the chance of false-negative results, the check also searches for base events outside of the specified date range. For example, you specify a date range of 29 May to 5 June. The check finds several verification events within the date range; however, Verification Event A was created on 29 May and includes base events that occurred on 28 May. To prevent the verification event from failing, the system will expand the search for base events beyond the specified dates.

> **NOTE**: The check process can take a long time if it includes large amounts of data. Therefore you should run the check during off-peak hours, and limit the date range to include only the data that you are interested in.

1. Select **Admin** > **Event Integrity**.

2. Specify the **Start Date** and **End Date** for the range of data that you want to check.
   The *Start Time* for the check corresponds to the time when you select **Run**. For example, 5:29 pm.

   > If the start and end dates encompass a time when the database is receiving events, it's possible to get a Missing Events notification. This occurs because the integrity check finds the verification events before ingesting their associated base events. We recommend that you avoid running the check against events currently being ingested by the database.

3. Select the timestamp type.

4. Click **Run**.

   If the **Run** button is disabled, a check is running currently. You can run one check at a time only. The check provides a **Status** update, as well as a showing when the check began and its specified date range.

5. (Optional) To cancel the check, click **Cancel**.

   Run as needed.

6. View the results.

View Event Integrity Check Results

To view the status and results, select **Admin** > **Event Integrity**.

The **Event Integrity Check** feature provides the following status and results:

- "View the Event Integrity Check Status" below
- "View Last Event Integrity Check Results Table" on the next page

View the Event Integrity Check Status

The **Event Integrity Check status** displays the date range from which the results are currently being checked, as well as the current status.

> If the system is busy, it might take a moment for the interface to indicate whether a check is running or has been canceled.

New

Indicates that you have never performed an Event Integrity Check; therefore, no results display.

In Progress

Indicates that an Event Integrity Check is running. If **Run** is disabled, a check is running currently. You can run one check at a time only.

Canceled

Indicates that you canceled an Event Integrity Check, and the system has completed the cancel task.

Canceling the check might take time to end the tasks that had been in progress.

Completed

Indicates that the Event Integrity Check has completed successfully. The results display in the results table.

Failed

Indicates that the Event Integrity Check failed to complete due to an error. The following table lists the failure categories and recommended remediation.

> Hover your cursor over the information icon next to "Failed" to display the corresponding error message.

For this error...	Which indicates that...	You might want to...
No Data Found	The check did not find verification events in the selected date range	Enable Event Integrity in the SmartConnector or Transformation Hub
Insufficient Disk Space	The system does not have enough disk space to run the check	Run the check when the system is less busy
Insufficient Resources	The system does not have enough memory to run the check	Run the check when the system is less busy
Event Integrity Running	An Event Integrity Check is already running	Wait for the current check to finish or cancel the check in progress
Other errors	Depends on the situation	Review the details of the error in the Search Engine log

View Last Event Integrity Check Results Table

The **Last Event Integrity Check Results table** displays the date range from the last check as well as the details of the last check, including the following information:

> The Event Integrity Check might display the last result from another user. The last result might also display after logging out.

Base events checked

Represents the number of base events referenced by verification events found in the specified date range.

Intact events

Represents the number of base events referenced by verification events that passed the Event Integrity check.

Missing events

Represents the number of base events referenced by verification events where missing base events exist.

> You might see this result when the integrity check finds the verification events before ingesting their associated base events. Try adjusting the start or end date for the check.

Tampering has been detected

Represents the number of base events referenced by verification events where the Event Integrity check failed to match the information provided by the verification events, such as

due to a change in the data in the base event.

Missing hashes, incorrect hashes, or base events being out of sequence usually indicates that the data has been deliberately changed in an attempt to hide a user's activities.

Duplicate base event IDs

Represents the number of base events referenced by verification events where the Event Integrity check failed because more than one base event has the same globally unique event ID. This situation results in the Event Integrity check of the referenced base events to fail.

Duplicate verification events IDs

Represents the number of base events referenced by verification events where the Event Integrity check failed because more than one verification event has a globally unique event ID.

Configure Data Collection to Support Event Integrity Checks

An Event Integrity Check can review both raw event data and the parsed fields within an event, depending on how you configure SmartConnectors and Transformation Hub, respectively.

Configure a SmartConnector to Generate a Verification Event

The Event Integrity Check can verify the raw data received from a SmartConnector. You must configure the connector to generate a verification event for batches of events. This configuration allows you to verify that the raw data in the database matches the event captured at the moment that it occurred in your environment.

To configure SmartConnectors, in a :

- Non-SaaS environment, see "Configuring a SmartConnector to Include a Verification Event for Raw Events" in the *Administrator's Guide for the ArcSight Platform*.
- SaaS environment, see "Configuring a SmartConnector to Include a Verification Event for Raw Events" in the *Quick Start for Administrators*.

Enable Transformation Hub to Generate Verification Events for Parsed Fields

The Event Integrity Check can verify the integrity of multiple parsed fields within an event that Transformation Hub received from a SmartConnector. For example, it verifies the *deviceProduct* and *sourceHostName* fields if they exist in an event. You might want to check these types of parsed fields as an alternative to verifying the raw event data. For example, your environment might not have the disk space, processing power, or network capacity to manage large amounts of raw event data forwarded from SmartConnectors.

You can configure this setting as you deploy Transformation Hub or at any time after deployment. For more information, in a non-SaaS environment, see "Enabling Transformation Hub to Generate Verification Events for Parsed Fields" in the *Administrator's Guide for the ArcSight Platform*. The SaaS environment does not currently support checks of parsed events.

Hunting for Threats and Vulnerabilities

Available only with ArcSight capabilities.

To help you hunt for undetected threats and vulnerabilities, the Reports Portal includes a set of built-in dashboards and reports. You can view this content based on the tactics and standards established by the Cloud Security Alliance and OWASP. Additional report and dashboards focus on fundamental security issues, such as monitoring firewalls and malware. For rapid access to your regular dashboards, you can configure the Reports Portal to display those dashboards by default.

Chapter 2: Understanding the Cloud Security Dashboards and Reports

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud**.

Cloud services providers are highly accessible, and the vast amount of data that they host makes them an attractive target for malicious users. To help you assess the security of services in the cloud, we provide dashboards and reports based on the industry-wide standards set by the Cloud Security Alliance (CSA). This alliance has identified the most significant security threats to the shared, on-demand nature of cloud computing. CSA refers to these issues as the **Treacherous 12**.

Reporting includes the following dashboards and reports, organized by the Treacherous 12 categories:

Category	Dashboards	Reports
"Abuse and Nefarious Use of Cloud Services" on page 202	DoS Originated from EC2 Instances EC2 Instances Communicating with Cryptocurrency Entity EC2 Instances Querying Domains Involved in Phishing Attacks EC2 Machines Involved in Suspicious Communication Email Spam Originated from EC2 Instances Nefarious Activity by an Unauthorized Individual from EC2 Suspicious Activity Reported by Microsoft Azure Trojans or Backdoors Installed on EC2 Instances	n/a
"Account Hijacking" on page 204	Account Hijacking Vulnerabilities Man in the Middle Attacks Phishing Attacks Principal Invoked an API Commonly used to Discover Information Associated with AWS account	Broken Authentication and Session Management

Category	Dashboards	Reports
"Advanced Persistent Threats" on page 205	Trojans or Backdoors Installed on EC2 Instances	n/a
"Data Breaches" on page 206	All Information Leakage Events Information Disclosure Vulnerabilities Organizational Information Leakage Personal Information Leakage	n/a
"Data Loss" on page 207	Amazon AWS Deletion Events	Amazon S3 Bucket Deletion Events Amazon VPC Deletion Events
"Denial of Service" on page 208	DoS Activity	n/a
"Insecure Interfaces and APIs" on page 209	n/a	Vulnerabilities on Interfaces and API
"Insufficient Due Diligence" on page 210	n/a	EC2 Machines Behavior Deviates from the Established Baseline Failed Technical Compliance Events
"Insufficient Identity Credential and Access Management " on page 211	n/a	AWS Account Password Policy Was Weakened Invalid or Expired Certificate Unsecured Password Events
"Malicious Insiders" on page 212	n/a	Nefarious Activity by an Unauthorized Individual
"System Vulnerabilities" on page 213	Vulnerability Overview	Cloud Related Vulnerabilities Critical Vulnerabilities Heartbleed Vulnerabilities Kernel Vulnerabilities Overflow Vulnerabilities Security Patch Missing Shellshock Vulnerabilities Spectre and Meltdown Vulnerabilities Vulnerabilities by Host
"Vulnerabilities on Shared Technologies" on page 215	n/a	"Vulnerabilities on Shared Technologies" on page 215

The cloud-based security dashboards and reports provide a view of events occurring in Amazon Web Service (AWS) and Azure, forwarded to the ArcSight Database from ArcSight ESM. Content in a dashboard depends on the widgets that it displays, as well as the dashboard's specified

time range. For example, some widgets summarize events by resource names and profile IDs, as well as by the event's severity.

Abuse and Nefarious Use of Cloud Services

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Malicious users can exploit poorly secured cloud service deployments, free cloud service trials, and fraudulent account sign-ups, which expose cloud computing models such as Iaas, PaaS, and SaaS. You might experience denial of service attacks, email spam and phishing campaigns, and brute-force computing attacks, or malicious individuals spoofing identities.

Some charts display data reported by Amazon GuardDuty, which is a threat detection service that continuously watches for malicious activity and unauthorized behavior.

Dashboards	Reports
DoS Originated from EC2 Instances	n/a
EC2 Instances Communicating with Cryptcurrency Entity	
EC2 Instances Querying Domains Involved in Phishing Attacks	
EC2 Machines Involved in Suspicious Communication	
Email Spam Originated from EC2 Instances	
Nefarious Activity by an Unauthorized Individual from EC2	
Suspicious Activity Reported by Microsoft Azure	
Trojans or Backdoors Installed on EC2 Instances	

DoS Originated from EC2 Instances
> Helps you identify denial of services activities that arise from EC2 (AWS Elastic Compute Cloud service) instances. The charts and table show events summarized by their Amazon resource name, severity, and GuardDuty.

EC2 Instances Communicating with Cryptocurrency Entity
> Displays EC2 instances that communicates with cryptocurrency IP addresses or domains.

EC2 Instances Querying Domains Involved in Phishing Attacks
> Lists the EC2 instances in which querying domains are involved in phishing attacks.

EC2 Machines Involved in Suspicious Communication
> Lists the EC2 machines that are involved in suspicious communication.

Email Spam Originated from EC2 Instances
> Identifies email spam that originates from EC2 instances.

Nefarious Activity by an Unauthorized Individual from EC2

Displays events that Amazon GuardDuty reports as nefarious activity by an unauthorized individual from EC2 machines. Amazon GuardDuty is a threat detection service that continuously watches for malicious activity and unauthorized behavior.

Suspicious Activity Reported by Microsoft Azure
Lists suspicious activity reported by Microsoft Azure.

Trojans or Backdoors Installed on EC2 Instances
Lists backdoors or trojans discovered on EC2 machines.

Account Hijacking

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > CSA > *The Treacherous 12.*

CSA identifies the hijacking of accounts and services as an ongoing, top threat. Malicious users might hijack accounts by phishing, fraud, and exploiting software vulnerabilities. In the cloud, the hijackers can eavesdrop on organizational activities, manipulate data, and redirect your clients.

Dashboards	Reports
Account Hijacking Vulnerabilities Man in the Middle Attacks Phishing Attacks Principal Invoked an API Commonly used to Discover Information Associated with AWS Account	Broken Authentication and Session Management

Account Hijacking Vulnerabilities
> Provides charts of the top 10 vulnerabilities and the number of vulnerabilities over time. This dashboard also includes a table of the vulnerabilities, so you can review the reporting vendor or device, agent severity, asset, and the asset's zone.

Man in the Middle Attack
> Provides charts that show man in the middle events by time, source address, destination address, source MAC address, and destination MAC address.

Phishing Attacks
> Provides charts that show the phishing attacks against the organization.

Principal Invoked an API Commonly used to Discover Information Associated with AWS account
> Provides charts that show the principals invoked by an API commonly used to discover information associated with AWS accounts.

Broken Authentication and Session Management
> Lists the events that might be associated with broken authentication (possibly hijacked credentials) and session management issues reported by vulnerability scanners in the organization.

Advanced Persistent Threats

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Advanced Persistent Threats (APTs) are a parasitical form of cyberattack that infiltrates systems to establish a foothold in the computing infrastructure of target companies from which they smuggle data and intellectual property.

Dashboards	Reports
Trojans or Backdoors installed on EC2 Instances	n/a

Trojans or Backdoors Installed on EC2 Instance
Provides charts showing backdoors or trojans discovered on EC2 (AWS Elastic Compute Cloud service) machines.

Data Breaches

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

While the risk of a data breach is not unique to the cloud, the CSA ranks it as a top concern for cloud customers. Sometimes the breach is the prime motivation of malicious users. However, breaches also result from mistakes made by individuals within the organization or poor security practices and software vulnerabilities.

To search for potential threats, use the following dashboards:

Dashboards	Reports
All Information Leakage Events Information Disclosure Vulnerabilities Organizational Information Leakage Personal Information Leakage	n/a

All Information Leakage Events
> Provides charts and a table that show the leakage events in the organization, including the top reported events, destination users, and assets.

Information Disclosure Vulnerabilities
> Provides charts and a table that show the disclosure vulnerabilities reported in the organization over time and by agent severity. You can also see the top 20 hosts, IP addresses, and signature ID events.

Organizational Information Leakage
> Provides charts and a table that show the leakage of organizational information. You can view the top 20 leakage events and signature IDs, as well as leakage over time and agent severity.

Personal Information Leakage
> Provides charts and a table that show the leakage of personal information. You can view the top reported, top 10 destination and source users, and leakage over time.

Data Loss

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

No organization wants to lose data, particularly to malicious individuals who might use the information in an adverse manner. Unfortunately, data stored in the cloud can also be deleted accidentally or as a result of a catastrophe.

Dashboards	Reports
Amazon AWS Deletion Events	Amazon S3 Bucket Deletion Events
	Amazon VPC Deletion Events

To assess the potential for data loss, use the following reports:

Amazon AWS Deletion Events
Provides charts and a table listing the number of deletion events by operations, day, source address, and source user.

Amazon S3 Bucket Deletion Events
Lists the deletion events that occur in Amazon S3 Buckets.

Amazon VPC Deletion Events
Lists the deletion events that occur in Amazon VPC.

Denial of Service

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Denial-of-service (DoS) attacks deliberately attempt to prevent users from accessing services, data, and applications. Use the **DoS Activity** dashboard to watch for potential service interruptions. You can view the top source and destination addresses, as well as events by day.

Dashboards	Reports
DoS Activity	

DoS Activity

Provides charts the top source and destination addresses, as well as events by day. This dashboard also is available in the Network Monitoring category of the **Foundation** reports.

Insecure Interfaces and APIs

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Users interact with cloud computing services through user interfaces (UIs) and application program interfaces (APIs), and the value-added services built on these services. APIs and UIs are generally the most exposed part of a system, perhaps the only asset with an IP address available outside the trusted organizational boundary. These assets will be the target of heavy attack.

Dashboards	Reports
n/a	Vulnerabilities on Interfaces and API

Vulnerabilities on Interfaces and API
   Reports the vulnerabilities found in your cloud-based interfaces and APIs.

Insufficient Due Diligence

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

The CSA states that it is essential to develop a good roadmap and checklist for due diligence when evaluating technologies and CSPs. Organizations should perform due diligence to mitigate the myriad risks associated with providing cloud services.

Dashboards	Reports
n/a	EC2 Machines Behavior Deviates from the Established Baseline
	Failed Technical Compliance Events

EC2 Machines Behavior Deviates from the Established Baseline
Details how the behavior of EC2 (AWS Elastic Compute Cloud) machines deviates from the established baseline.

Failed Technical Compliance Events
Lists the failed technical compliance events.

Insufficient Identity Credential and Access Management

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Malicious users can infiltrate and cause data breaches based on poor authentication methods and weak password policies.

Dashboards	Reports
n/a	AWS Account Password Policy Was Weakened
	Invalid or Expired Certificate
	Unsecured Password Events

AWS Account Password Policy Was Weakened
> Lists events associated with weakened AWS account password policy.

Invalid or Expired Certificate
> Lists events associated with invalid or expired certificates.

Unsecured Password Events
> Lists events associated with unsecured passwords.

Malicious Insiders

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Individuals within an organization, such as system administrators or disgruntled colleagues, might access sensitive information for malicious intent. Most organizations use controls to limit risk from malicious insiders, such as controlling encryption keys and monitoring or auditing the activities of specific users.

Dashboards	Reports
n/a	Nefarious Activity by an Unauthorized Individual

Nefarious Activity by an Unauthorized Individual

Lists events that Amazon GuardDuty reports as nefarious activity by an unauthorized individual from EC2 (AWS Elastic Compute Cloud) machines. Amazon GuardDuty is a threat detection service that continuously watches for malicious activity and unauthorized behavior.

System Vulnerabilities

Select > **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **System Vulnerabilities**.

Most computer systems have programs, services, and operating systems that are vulnerable to exploitation. According to the CSA, vulnerabilities within the components of the operating system – kernel, system libraries and application tools – put the security of all services and data at significant risk.

Dashboards	Reports
Vulnerability Overview	Cloud Related Vulnerabilities
	Critical Vulnerabilities
	Heartbleed Vulnerabilities
	Kernel Vulnerabilities
	Overflow Vulnerabilities
	Security Patch Missing
	Shellshock Vulnerabilities
	Spectre and Meltdown Vulnerabilities
	Vulnerabilities by Host

Cloud Related Vulnerabilities

Lists all events associated with vulnerabilities known to affect AWS and Azure.

Critical Vulnerabilities

Lists all events that have a High or Very High severity, based on CVE and CVSS data.

Heartbleed Vulnerabilities

Lists all events associated with the heartbleed bug, which is a system vulnerability in the OpenSSL cryptographic software library. This weakness allows malicious users to steal the information protected, under normal conditions, by the SSL/TLS encryption used to secure the Internet. A Heartbleed attack works by tricking servers into leaking information stored in their memory. Attackers can also get access to a server's private encryption key, allowing the attacker to unscramble any private messages sent to the server and even impersonate the server.

Kernel Vulnerabilities

Lists all events associated with kernel vulnerabilities. For example, the vulnerability in the Linux Kernel netfilter/xt_TCPMSS, which could allow remote hackers to carry out a denial of service attack.

Overflow Vulnerabilities

Lists all events associated with buffer overflows. When a buffer receives more data than it can handle, the data can overflow to other storage locations. Overflows can cause system crashes or create an exploitable vulnerability.

Security Patch Missing

Reports the hosts that do not have the security patches needed to resolve known vulnerabilities.

ShellShock Vulnerabilities

Reports the hosts vulnerable to a ShellShock attack. In a ShellShock attack, the Unix shell Bash could execute arbitrary commands and allow unauthorized access to services, such as web servers, that use Bash to process requests.

Spectre and Meltdown Vulnerabilities

Reports the hosts vulnerable to Meltdown and Spectre attacks, which exploit critical vulnerabilities in modern processors. Meltdown breaks the fundamental isolation between user applications and the operating system, allowing a program to access the memory and data of other programs and the operating system. Spectre attacks break the isolation between applications, allowing programs to leak information to each other. These exploitations do not leave any traces in traditional log files.

Vulnerability Overview

Provides a dashboard view of the vulnerabilities found in the organization.

Vulnerabilities by Host

Lists all vulnerabilities detected on the specified hosts.

Vulnerabilities on Shared Technologies

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Cloud** > **CSA** > **The Treacherous 12**.

Some technologies that form the infrastructure for the cloud-based services started as on-premises capabilities, and thus might not have been designed to share its resources in multi-tenancy or multicustomer environments. For example, an application might not have initially been expected to support multi-factor authentication or its database designed to partition data by tenant.

Dashboards	Reports
n/a	Vulnerabilities on Shared Technologies

Vulnerabilities on Shared Technologies
 Lists the vulnerable technologies that a malicious user might exploit.

Understanding the Foundation Dashboards and Reports

Available only with ArcSight capabilities.

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

Reporting includes the following dashboards and reports, organized by the following foundational categories:

Category	Dashboards	Reports
"Entity Monitoring" on page 218	Account Management Overview Failed Logins Overview Successful Login Overview	All Logins by Hostname Failed Logins Summary Login Activity by User
"Events Overview" on page 220	Least Common Events Most Common Events Most Common Events by Severity Reporting Devices	n/a
"Hosts Monitoring" on page 221	Host Profile Overview	Anti-Virus Activity Anti-Virus Stopped or Paused Audit Log Cleared Failed Anti-Virus Updates Summary Operating Systems Errors and Warnings Services Shutdown Services Started
"Malware Monitoring" on page 223	Malware Overview	Reported Malware by Host Worm Infected Systems

Category	Dashboards	Reports
"Network Monitoring" on page 224	Attacks and Suspicious Activity Overview DGA Overview DoS Activity Email Attacks IDS Events Man in the Middle Attacks Reconnaissance Activity Traffic Anomaly Overview VPN Activities Overview	Exploit Attempts Detected by IDS Network Device Configuration Changes
"Perimeter Monitoring" on page 226	Firewall Blocked Events Firewall Traffic Overview	Firewall Configuration Changes Firewall Blocked Traffic by Destination Address
"Vulnerability Monitoring" on page 227	n/a	High Risk Vulnerabilities by Host SSL Vulnerabilities Vulnerability Overview Vulnerabilities by Host XSRF Vulnerabilities XSS Vulnerabilities

Entity Monitoring

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

To prevent brute force attacks or denial-of-service attacks, you could track login activities in your environment. A malicious user might attempt to guess another user's password by repeatedly attempting to log in to the same account. You can track this behavior by observing failed login attempts. You might also watch for users who attempt to log in to multiple devices and hosts. Malicious users might also create, modify, and delete accounts to gain unauthorized access or let them execute harmful code.

To monitor account activity, use the following dashboards and reports:

Dashboards	Reports
Account Management Overview	All Logins by Hostname
Failed Logins Overview	Failed Logins Summary
Successful Login Overview	Login Activity by User

Account Management Overview

Provides charts and a table to help you identify users who are creating and deleting the most accounts. You also can track which hosts have had the largest number of accounts modified or deleted.

All Logins by Hostname

Reports the number of login attempts over time, including the outcome, for the specified hosts.

You must specify one IP address.

Failed Logins Overview

Provides an overview, in charts and a table, of the hosts and users with the highest number of failed logins. You can also view the number of failed logins over time, by reporting device, or source address.

Failed Logins Summary

Reports the number of failed logins over time. The table includes the user, source address, target host, and number of failed attempts.

Login Activity by User

Reports the number of times that the specified users have attempted to log in to a host. The table indicates whether the attempt is successful.

You must specify one user by Destination UserName.

Successful Login Overview

Provides an overview, in charts and a table, of users with the highest number of successful logins. You can review the relationship between the users and the hosts to which they successfully log in.

Events Overview

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

To identify threats in your environment, you might want to have an overview of the events that occur the most often or affect the most devices and hosts. You could also watch for events that rarely occur to check for unusual activity.

To monitor event activity, use the following dashboards:

Dashboards	Reports
Least Common Events	n/a
Most Common Events	
Most Common Events by Severity	
Reporting Devices	

Least Common Events

Provides charts and a table to help you identify the events that have the fewest reported occurrences. You can view the results by vendor, such as Amazon, or product, such as Microsoft Windows.

Most Common Events

Provides charts and a table to help you identify the common events that affect your environment by vendor, such as Amazon, or product, such as Microsoft Windows.

Most Common Events by Severity

Provides a table to help you track the events by count and severity.

Reporting Devices

Provides charts and a table to help you identify the hosts and devices with the most reported security events. You can view charts summarizing the most common severity of the events; top 20 events by vendor such as Microsoft or McAfee; top 20 events types of events, such as stopped services, and the top 20 events by class ID, such as a CVE.

Hosts Monitoring

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

In general, you should consistently monitor host-based events that indicate unauthorized activities. For example, a malicious user or program might start and stop host services and anti-virus programs. Additionally, they might clear the audit log to hide their actions on a host.

To monitor unusual activity that affects hosts, use the following reports:

Dashboards	Reports
Host Profile Overview	Anti-virus Activity
	Anti-virus Stopped or Paused
	Audit Log Cleared Events
	Failed Anti-virus Updates Summary
	Operating System Errors and Warnings
	Services Shutdown
	Services Started

Host Profile Overview

> Displays the activity on a specific host. This is a drill-down dashboard that can be opened from IP addresses or host names located in a table, bar, chart, or entered in the dashboard search bar. The dashboard requires a valid IP address or host name to show information.

Anti-virus Activity

> Reports the volume of activity by reporting anti-virus service. The table provides results by event name, count, affected host, and outcome.

Anti-virus Stopped or Paused

> Reports the top IP addresses where an anti-virus service has been stopped or paused. The table provides results by host, service name, and number of events.

Audit Log Cleared

> Reports the number of times that the audit log has been cleared by user, host, and date.

Failed Anti-virus Updates Summary

> Reports the number of failures in updating anti-virus software by date and host.

Operating Systems Errors and Warnings

> Reports the top system errors and warnings by host. You could identify issues associated with specific errors or warnings, such as privileged objects and users, password changes, and login failures. Alternatively, you could sort the table by the reported hosts to review the types of issues affecting each host.

Services Shutdown

Reports the top 10 services that have been shut down in your environment. The table provides a summary of all services, including the associated hosts.

Services Started

Reports the top 10 services that have been started in your environment. The table provides a summary of all services started, including the associated hosts.

Malware Monitoring

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

Malware, or malicious software, represents all the variations of programs designed to damage computers, servers, clients, devices, applications, and networks.

To monitor unusual activity that affects hosts, use the following reports:

Dashboards	Reports
Malware Overview	Reported Malware by Host
	Worm Infected Systems

Malware Overview

Provides charts and a table to help you identify the malware affecting your enterprise and the top 10 infected hosts. You can also view the malware events reported over time.

Reported Malware by Host

Lists the malware found on the specified hosts.

You must specify one host.

Worm Infected Systems

Lists the hosts infected by worms, and provides a chart that shows the malware by count found in your enterprise.

Network Monitoring

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

The traffic exchanged between devices and servers tells you a lot about your network. By monitoring network traffic, you can identify cyber attacks and network events that could affect your enterprise. For example, malicious users might find a way to intercept communications to generate a man-inthe-middle attack or change the configuration of devices to gain unauthorized access. In both cases, the attack is the beginning of further intrusions. Also, a system infected by malware can be instructed generate a large volume of domains, thus causing increased traffic.

To monitor network activity, use the following dashboards and reports:

Dashboards	Reports
Attack and Suspicious Activity	Exploit Attempts Detected by IDS
DGA Overview	Network Device Configuration Changes
DoS Activity	
Email Attacks	
IDS Events	
Man in the Middle Atacks	
Reconnaissance Activity	
Traffic Anomaly Overview	
VPN Activities Overview	

Attack and Suspicious Activity Overview

Provides charts and a table to help you identify the top attackers, targets, and events over time.

DGA Overview

Provides charts and a table to help you watch for domain generation algorithms (DGAs). You can identify the IP addresses generating the most DGA domains or the unique domains that the largest number of hosts attempt to connect with. You can also check for the hosts that are transmitting the largest amount of data.

DoS Activity

Provides charts and a table for you to identify denial-of-service events. You can view the number of events per day, as well as the top source and destination addresses.

This dashboard also is available in the Denial of Service category of the Cloud reports.

Email Attacks

Provides charts and a table that describe the email attacks detected in your enterprise. You can view the top events or target users, as well as the destination and source addresses.

Exploit Attempts Detected by IDS

Shows the top 10 exploit attempts reported by the intrusion detection systems (IDS) in your enterprise. In the table, you can sort the events by count or severity.

IDS Events

Provides a chart and table showing all events reported by the IDSs in your enterprise.

Man in the Middle Atacks

Provides charts and a table to help you catch potential man-in-the-middle (MitM) attacks. You can view events over time, by source and destination address including MAC addresses, and the top MitM events.

During a MitM attack, the malicious user intercepts communications between two parties either to secretly eavesdrop or modify traffic traveling between the two.

Network Device Configuration Changes

Reports the top 10 devices whose configurations have changed, as well as the top 10 events causing configuration changes.

Reconnaissance Activity

Provides charts and a table to help you watch for active reconnaissance attacks. You can view identify the top sources of recon activity, as well as the primary destinations for these attacks. Review the pie charts to identify the main types of events and affected zones.

Active reconnaissance is a type of computer attack in which an intruder engages with the targeted system to gather information about vulnerabilities. Malicious users might use tools like ping or traceroute to perform recon through automated scanning or manual testing.

Traffic Anomaly Overview

Provides charts to help you identify anomalies in network traffic. You can view the top source and destination address, events, and activity over time.

VPN Activities Overview

Provides charts and a table for you to monitor VPN activity, such as the top users who access the VPN. You can view the VPN activities per day, as well as review the top source and destination addresses.

Perimeter Monitoring

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

The perimeters of an enterprise's network handle a great deal of traffic, causing system administrators to face an ever-increasing need to allow fast, efficient flow of traffic while also keeping the network secure. If you pro-actively monitor the firewalls in your enterprise, you can identify problems at an early stage and prevent network attacks. Malicious users often exploit loopholes in your firewall rules, particularly any old or unused rules. Network traffic also can be vulnerable to unencrypted data.

To monitor your network's perimeter, use the following dashboards and reports:

Dashboards	Reports
Firewall Blocked Events	Firewall Configuration Changes
Firewall Traffic Overview	Firewall Blocked Traffic by Destination Address

To monitor your network's perimeter, use the following dashboards and reports:

Firewall Blocked Events
Provides charts and a table for you to monitor the events that your firewalls have blocked, such as the bytes in and out for all blocked events. You can view the top events blocked per device, application protocol, source address, or destination address.

Firewall Blocked Traffic by Destination Address
Lists the top 10 firewall traffic events that have been blocked from reaching the specified hosts.

You must specify one IP address.

Firewall Configuration Changes
Lists the top 10 changes to the firewall configuration by host.

Firewall Traffic Overview
Provides charts and a table for you to monitor traffic through your firewalls, such as the bytes in and out by accepted and denied traffic. You can view the top reporting devices and destination addresses, as well as the outcomes of port usage over time. The table lists the Port, transport protocol, application protocol, and number of events reported by firewalls.

Vulnerability Monitoring

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **Foundation**.

Many of the components within a web application, such as the libraries and modules, run with the same privileges as the application itself. Applications and APIs using components with known vulnerabilities can undermine application defenses and enable various attacks and impacts. For example, malicious users can exploit a known in SSL with the Heartbleed Bug. Web site and web applications can be vulnerable to cross-site scripting (XSS) and cross-site request forgery (XSRF) attacks. In an XSRF attack, also known as a one-click attack or session riding, a malicious user submits unauthorized commands to a web application from a user account that the application trusts.

High-risk vulnerabilities represent those that are relatively easy for attackers to exploit and gain control over system components. Many high-risk vulnerabilities can temporarily or permanently disrupt enterprise operations.

To check whether your enterprise has vulnerabilities, use the following dashboard and reports:

Dashboards	Reports
Vulnerability Overview	High Risk Vulnerabilities by Host
	SSL Vulnerabilities
	Vulnerabilities by Host
	XSRF Vulnerabilities
	XSS Vulnerabilities

High Risk Vulnerabilities by Host

Lists all high-risk vulnerabilities found on the specified hosts.

You must specify one host by `Destination Host`.

SSL Vulnerabilities

Lists the hosts reported to have the most SSL vulnerabilities.

Vulnerability Overview

Provides charts and a table to help you track the vulnerabilities reported in your enterprise.

Vulnerabilities by Host

Lists all vulnerabilities found on the specified hosts.

You must specify one IP address.

XSRF Vulnerabilities

Lists the top 10 hosts that are vulnerable to a cross-site request forgery (XSRF or CSRF) attack.

XSS Vulnerabilities

Lists the top 10 hosts that are vulnerable to cross-site scripting (XSS) attacks.

Chapter 2: Understanding the OWASP Security Dashboards and Reports

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP**.

We provide dashboards and reports based on the industry-wide standards set by the Open Web Application Security Project®. OWASP is a nonprofit foundation that works to improve the security of software. The organization has established a list of the Top 10 security risks to web applications, focusing on the most critical threats to the shared, on-demand nature of webbased applications.

Reporting includes the following dashboards and reports, organized according to **OWASP's Top 10** risk categories:

Category	Dashboards	Reports
"Broken Access Control" on page 231	N/A	"Broken Access Control" on page 231
Cryptographic Failure	Information Leaks Overview	Organizational Records Information Leaks Personal Information Leaks
"Identification and Authentication Failures" on page 238	N/A	Broken Authentication and Session Management
"Injections" on page 233	Injection Vulnerabilities Overview XSS Vulnerabilities	Command Injections on HTTP Request Cross Site Scripting Injection Vulnerabilities SQL Injection
"Software and Data Integrity Failures" on page 239	Deserialization Flaws Overview	Deserialization Flaws
"Security Logging and Monitoring Failures" on page 240	Attacks and Suspicious Activity Overview Failed Logins Overview Login Activty Overview Security Log is Full	All Logins by Hostname Audit Log Cleared Failed Logins Summary Operating System Errors and Warnings

Category	Dashboards	Reports
"Security Misconfiguration" on page 235	Misconfiguration Events Overview Missing Security Patches Overview XML Vulnerabilities Overview	Security Patch Missing XML Vulnerabilities
Server-Side Request Forgery	N/A	Server-Side Request Forgery
"Vulnerable and Outdated Components" on page 237	SSH Vulnerabilities Overview Vulnerability Overview	SSH Vulnerabilities Summary SSL Vulnerabilities

Broken Access Control

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 1 - Broken Access Control**.

Some enterprises fail to enforce access controls that restrict what authenticated users are allowed to do. By exploiting vulnerabilities in access controls, a malicious user might retrieve sensitive files, gain access other user's accounts, change access rights, and misuse data.

Dashboards	Reports
n/a	Broken Access Control

Broken Access Control
  Lists vulnerable hosts by severity over time.

Cryptographic Failures

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 2 - Cryptographic Failures**.

Most enterprises store sensitive data that needs to be protected, such as personal information, customer and organizational financial data, healthcare records, or intellectual property. Web applications and APIs might inadvertently expose sensitive data by not having enough protections such as encryption at rest or in transit, or when exchanging data with the browser. Malicious users could use the data for credit card fraud, identity theft, and other crimes.

Dashboards	Reports
Information Leaks Overview	Organizational Records Information Leaks
	Personal Information Leaks

Information Leaks Overview
Provides charts and a table to help you identify the most reported systems, types of leaks, and leakage events that occur over time. You can identify the top reported users and view leaks by category.

Organizational Records Information Leaks
Lists the top leakage events that affect organizational records.

Personal Information Leaks
Lists the top leakage events that affect personal records by Destination UserName.

Injections

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 3 - Injections**.

Injection vulnerabilities, or flaws, allow malicious users to inject code in other systems, especially interpreters, by using vulnerable applications. For example, in a SQL, NoSQL, OS or LDAP injection attack, someone sends untrusted data to an interpreter as part of a command or query to trick the interpreter into executing hostile commands or accessing data without appropriate authorization. Usually, these flaws result from insufficient validation of data input or the failure to filter or sanitize the input.

Dashboards	Reports
Injection Vulnerabilities Overview	Command Injections on HTTP Request
XSS Vulnerabilities	Cross Site Scripting
	Injection Vulnerabilities
	SQL Injection

Command Injections on HTTP Request

Lists the highest number of events associated with command injections in an HTTP request, by the requested URL. This report includes a chart to help you identify the relationship between the IP addresses of the attacker and the target.

In a command injection attack that exploits an HTTP request, malicious users execute arbitrary commands on the host operating system via a vulnerable application. For example, the web application passes unsafe data supplied by the user to a system shell.

Cross Site Scripting

Lists events associated with XSS vulnerabilities.

Injection Vulnerabilities

Lists the hosts with the most injection vulnerabilities over time.

Injection Vulnerabilities Overview

Provides charts and a table to help you identify the systems affected by injection vulnerabilities, as well as view the top reported vulnerabilities by agent severity, risk, and over time.

SQL Injection

Lists the systems with the highest number of SQL injection vulnerabilities.

In a SQL injection attack, a malicious user can interfere with the queries that an application makes to its database. The user could view delete, or modify data not usually available for

retrieval. A malicious user could also use SQL injections to start a denial-of-service attack or compromise other services, servers, or infrastructure.

Vulnerabilities associated with **cross-site scripting (XSS)** enable malicious users to inject code in legitimate web pages or applications that executes harmful scripts in the user's web browser when the browser parses data. The scripts might hijack user sessions, deface web sites, or redirect users to harmful sites. A web application or web page becomes vulnerable when it includes untrusted data; data without proper validation or escaping; or data supplied by users through an API that can create HTML or Java-script. XSS attacks tend to occur in forums, message boards, and web pages that allow comments. Malicious users can execute XSS attacks in VPSCript, ActiveX, Flash, and CSS. However, this type of injection attack most commonly occurs in Java Script.

XSS Vulnerabilities

Provides charts and a table so you can review potential XSS vulnerabilities in your environment by vulnerability type or the top vulnerable hosts.

To get a list of the top 10 hosts vulnerable to cross-site scripting attacks, run the XSS Vulnerabilities report.

Security Misconfiguration

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 5 - Security Misconfiguration**.

In general, the most common vulnerability in your environment is mis-configured operating systems, frameworks, libraries, and applications. Mis-configurations include missing security patches or updates, incomplete or ad hoc configurations, use of insecure default configurations, poorly configured HTTP headers, and error messages that contain sensitive information.

Dashboards	Reports
Misconfiguration Events Overview	Security Patch Missing
Missing Security Patches Overview	XML Vulnerabilities
XML Vulnerabilities Overview	

Misconfiguration Events Overview
Provides an overview of the mis-configured events reported in your environment. The charts show the top mis-configured systems, the top misconfiguration events, an indicator of the risk associated with the reported misconfiguration events, events by agent severity, and misconfiguration events over time. The table provides additional information, such as the associated vulnerability.

Missing Security Patches Overview
Provides charts and a table to help you identify the top machines that fail to have all relevant security patches, as well as the security patches most reported as not having been applied. You can review the missing patch reports over time, by agent severity, and by risk indicator.

Security Patch Missing
Lists the security patches that have not been applied, as reported by vulnerability scanners in your environment.

Older or mis-configured XML processors use XML documents to evaluate external entity references, and can inadvertently process harmful XML input. Malicious users the XML processor's to reveal internal content such as files, file shares, and port scans, as well as execute remote code and denial-of-service attacks.

XML Vulnerabilities
Lists the hosts with the most XML vulnerabilites.

XML Vulnerabilities Overview>

Provides charts and a table to help you identify the systems with the most XML vulneraibilitie as well as the most reported vulnerabilites. You can review the vulnerabilities by severity and risk indicator.

Vulnerable and Outdated Components

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 6 - Vulnerable and Outdated Components**.

Many of the components within a web application, such as the libraries and modules, run with the same privileges as the application itself. Applications and APIs using components with known vulnerabilities can undermine application defenses and enable various attacks and impacts. Malicious users can exploit vulnerabilities in SSH and SSL. For example, the Heartbleed Bug is a known SSL vulnerability. Your enterprise might have large numbers of SSH keys because end users can create new SSH keys (credentials) or even duplicate them without oversight, unlike certificates or passwords. A malicious user can gain long-term access to your resources by taking advantage of SSH keys that have been left unaccounted for.

Dashboards	Reports
SSH Vulnerabilities Overview	SSH Vulnerabilities Summary
Vulnerability Overview	SSL Vulnerabilities

SSH Vulnerabilities Overview

Provides charts and a table that show hosts with the most SSH vulnerabilities and the most reported vulnerabilities. You can review these vulnerabilities over time, by agent severity, and by risk indicator.

SSH Vulnerabilities Summary

Lists the hosts reported to have the most SSH vulnerabilities.

SSL Vulnerabilities

Lists the hosts reported to have the most SSL vulnerabilities.
This report also is available in the Vulnerability Monitoring category of the **Foundation** reports.

Vulnerability Overview

Provides charts and a table that show the top signature IDs for the anti-virus programs that have failed to update, as well as the hosts most likely to be vulnerable. You can review these vulnerabilities over time and by agent severity.

Identification and Authentication Failures

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 7 - Identification and Authentication Failures**.

Some enterprises mis-configure or fail to enable the authentication and session management functions of applications and web sites. When this occurs, a malicious user could compromise passwords, keys, and session tokens.

Dashboards	Reports
n/a	Broken Authentication and Session Management

Broken Authentication and Session Management

Reports the top 20 hosts with the most reports of broken authentication and system management. The table lists the IP address, host name, ID of the device event class, and the number of reported events. This report also is available in the Account Hijacking category of the **Cloud** reports.

Software and Data Integrity Failures

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 8 - Software and Data Integrity Failures**.

Untrusted, or insecure, deserialization allows malicious users to use untrusted data to abuse the logic of an application, initiate a denial-of-service or injection attacks, or execute harmful code when the data is deserialized. The user could even replace a serialized object with objects of a different class. Deserialization is a common process where the web site or application takes data from a file, stream, or network and rebuilds it into an object. The serialized objects might be used in JSON, XML, or YAML.

Dashboards	Reports
Deserialization Flaws Overview	Deserialization Flaws

Deserialization Flaws
> Lists the hosts with most deserialization flaws.

Deserialization Flaws Overview
> Provides charts and a table to help you identify the top hosts, deserialization flaws, and flaws found over time. You can view the flaws by agent severity and risk indicator.

Security Logging and Monitoring Failures

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 9 - Security Logging and Monitoring Failures**.

According to OWASP, insufficient logging and monitoring, coupled with missing or ineffective integration with incident response, allows malicious users to further attack systems; maintain persistence; pivot to more systems; and tamper, extract, or destroy data. Most major incidents start with an exploitation of the vulnerabilities in logging and monitoring. Yet, most organizations fail to discover the breach until several months have passed.

To help you detect potential breaches as soon as possible, use the following reports and dashboards:

Dashboards	Reports
Attacks and Suspicious Activity Overview	All Logins by Hostname
Failed Logins Overview	Audit Log Cleared
Login Activity Overview	Failed Logins Summary
Security Log is Full	Operating System Errors and Warnings

All Logins by Hostname
> Lists all logins that have occurred on the specified host.

Attacks and Suspicious Activity Overview
> Provides charts and a table to help you identify the top attackers, targets, and events over time.

Audit Log Cleared
> Lists all the Audit Clear events that have occurred in the organization.

Failed Logins Overview
> Provides charts and a table showing failed logins by time, users, hosts, reporting devices, and attacker address.

Failed Logins Summary
> Lists the failed login events that have occurred in your environment.

Login Activity Overview
> Provides charts and a table showing the outcome of login activity, including successful logins. You can view activity by machine or user, as well as a chart showing the relationship between users and systems to which they log in.

Operating System Errors and Warnings

Provides charts and a table that report the operating systems errors and warnings in the organization.

Security Log is Full

Provides charts and a table to help you identify the hosts where the security log is full.

Server-Side Request Forgery

Select **Reports** > **Portal** > **Repository** > **Standard Content** > **OWASP** > **A 10 - Server-Side Request Forgery**.

When server-side web applications fetch remote resources without checking the user-supplied URL, attackers may force the application to send a crafted request to an unexpected destination including those protected by a firewall or VPN.

Dashboards	Reports
n/a	Server-Side Request Forgery

Server-Side Request Forgery

Lists any Server-Side Request Forgery (SSRF) events that occur in your environment. This report provides details including: relevant hostnames, event IDs captured, the product that captured the events, and the number of times the event occurred. Additionally, this report has a chart for a visual representation of the information.

Responding to Threats

Select **RESPOND**.

ArcSight SOAR helps you respond quickly to threat detection and remediation by providing multiple forms of automation, analyst augmentation, collaborative investigation and response. With SOAR fully integrated in the Fusion capability, the system listens to alerts, forwarded from different components to identify a threat possibility. For Fusion to receive an alert, you must ensure that SOAR is integrated with other components.

Creating User Credentials for Integration

Select **RESPOND > Configuration > Credential**.

Fusion in the ArcSight Platform User's Guide listens to alerts, forwarded from different components to identify a threat possibility. For Fusion to receive an alert, you must ensure that SOAR is integrated with other components. The **Credentials** tab allows you to create user credentials to interact with other components during the integration procedure.

The Credentials tab displays a list of user credentials. You can view the credential names, the last modification date and the name of the modifier.

- Searching a Credential
- Creating a User Credential
- Editing and Deleting a User Credential

Searching a Credential

You can search a specific user credential, through the **Search** option. Click the button next to search, allows you to view search results based on **ID, Credential Name** and **Last Modified By**.

Creating a User Credential

Click the **+Create Credential** button to create a new user credential. In the **Credential Editor** window, specify the details for following fields:

Type
Select <Internal Credential, External Condition>.
Internal Credentials are stored in SOAR's database table. External Credentials are stored in integrations, such as Cyberark Central Credential Provider.

Name
Specify a name for the credential set.

The name that you create here is displayed in the **Credentials** field during alert source and integration configuration. You must select this name to ensure that SOAR communicates with other components through this name identity.

Username
Specify a username for the credential set.

Password
Specify a password for the credential set.

Private Key
Specify a private key for the credential set, if needed.

Editing and Deleting a User Credential

To edit an existing user credential, complete the following steps:

1. Click **Edit** in the **Actions** column.

2. In the **Credential Editor** window, specify the values as per your requirement.

3. Click **Save**.

4. (Optional) To delete an existing credential, click **Delete**.

> You cannot delete a user credential that is used in the integration with other components.

Configuring an Alert Source

Select **RESPOND > Configuration > Alert Source**.

To ensure seamless reception of the correlated alerts, you must configure an alert source. The **Alert Source** tab allows you to create new alert source configurations, displays a list of existing alert source configuration and options to modify the existing alert source configuration.

- Creating an Alert Source Configuration
- Editing and Deleting an Alert Source Configuration
- Configuring SOAR as an Alert Source

Creating an Alert Source Configuration

Click **+ Create Alert Source Configuration** button to create a new alert source configuration. In the **Alert Source ConfigurationEditor** window, specify the details for following fields:

> You might see differences in the fields of this editor for some of the alert source types (as you select it from the Type combo box list).

Name
Represents the name of the alert source.

Type
Represents the type of the alert source. It could be one of the alert source types.

Address
Represents the IP address of the alert source to which SOAR connects when it wants to get data.

Key
Represents a unique, auto-generated key that is used as a shared token to make sure the remote IP addresses ("Allowed IP addresses") are correct. The value of this field must be included in the messages coming from those remote IP addresses.

Allowed IP addresses
Indicates that any alert coming from the specified IP addresses will be processed and others will be discarded. For most alert source types, SOAR opens a TCP port (or a web service API endpoint) and waits for some alert sources to connect. This field along with the "Key" field is to improve your system's security. The combination of these two fields prevents a potential attacker from feeding your system with fake events and causing damages.

Alert Severity

Represents the severity of alert sources. Define the severities according to the priorities of tickets produced by the alert source. Use the **Add** button to create each severity. While adding the severities, you can specify the default severity by selecting the checkbox under the **Default** column.

Configuration Content

Indicates the default configuration definitions for some type of alert sources, such as IBM Security QRadar but it is not required for many alert sources. It depends on which alert source you are trying to interact with. If there are some required data for the alert source configuration, this area shows a template and ask you to edit it if needed.

Credential

Represents the credentials defined on the system to be used for the alert source.

Show alert parameters by default

Shows the default alert parameters defined for the selected device type on the system.

Trust Invalid SSL Certificates

Instructs SOAR to connect anyways to an alert source ignoring warnings for untrusted SSL certificates. You might have installed alert sources with self-signed SSL certificates, which SOAR does not trust and deny connecting by default. Therefore, if you do not select this checkbox, SOAR still gets the brief alert, but cannot get more details on the alert.

You can click **Test** to verify if the configuration is correct.

Editing and Deleting an Alert Source Configuration

You can edit an existing alert source configuration by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Alert Source Configuration Editor** window is displayed. Specify the values in the window as per your requirement and click **Save** to modify.

You can delete an existing alert source configuration by clicking the **Delete** button under the **Actions** column.

For detailed information see the related *Integration Guides*.

Configuring SOAR as an Alert Source

ArcSight SOAR creates internal alerts for some cases, such as when an action is failed permanently or an integration becomes unavailable because its firewall is not reachable. These internal alerts are generated for the event types including : action and rollback failures, auto-

enrichment failures, when an integration becomes offline/online, breach of ticket resolution/first response SLAs, and custom/arbitrary alerts created by playbooks.

Integrating SOAR with Other Components

Select **RESPOND > Configuration > Integrations**.

You can configure SOAR to integrate with other platforms and components to receive alerts. This procedure ensures streamlining the alert inflow and powers automation. The **Integrations** tab allows you to create, manage and configure security integrations and platforms. The Integrations page lists integrations configured previously along with their action and rollback queue sizes, and their availability statuses.

- Searching an Integration
- Creating an Integration
- Editing and Deleting an Integration
- Testing Integration
- Flushing Queues

Searching an Integration

You can search a specific integration, through the **Search** option. Click the button next to search, to view search results based on **ID, Name, Type, Address, Availability, Last Modified By, Modification Date, Action Queue Size, Rollback Queue Size** and **Actions** filters.

Creating an Integration

Click the **Create Integration** button to create an integration. In the **Integration Editor** window, specify the details for following fields:

Name
Name of the integration.

Type
Type of the integration.

IP Address
IP address of the integration.

Configuration

Depending on the integration type, you might select and enter various configuration commands on the black window. See the below Changing Integration Configuration section for details.

Credential

Credentials to be used to connect this integration. Credentials are defined in **Credentials** menu.

Trust Invalid SSL Certificates

Determines whether SOAR connects to an integration and ignores warnings for untrusted SSL certificates.

Require Approval Form

Identifies the users that need to approve action items before executing it for integrations.

Notify

Identifies the users that will be notified of actions done.

Tags

Used to group integrations. This allows creating actions on a number of integrations having the same tag. You might want to create an action for all integrations that have a specified tag such as "block offender IP address on all firewalls that are used to manage WiFi networks".

You might prefer to specify some more parameters for some specific integrations. Select the **Show Additional Parameters** checkbox located at the very bottom of the **Integration Editor** to the additional configuration.

Maintenance

Maintenance is supported by all integrations to which SOAR connects using the SSH protocol. It is essentially a generic SSH integration action script. It is best used in conjunction with Check Point Firewall integration for activating or installing a previously saved but not activated firewall policy. You can select a maintenance frequency or type your own cron job (for a scheduled maintenance) by selecting the **Custom Cron Value** option in the combobox.

Host Key

SSH public key of the remote integration. It is only used for integrations connected with SSH. If an SSH key is provided, then it will be validated using the specified key. This check is required to prevent man-in-the-middle attacks.

Batch Size

SOAR can send multiple action queue items to the integrations in a single connection. This field specifies the maximum number of action queue items that will be sent in each execution. For example, if you provided **Batch Size** as **10** and there are 25 action queue

items waiting for that integration, then SOAR will send these items in 3 separate execution (10 + 10 + 5). Its default value is 1. This is a feature to avoid causing excessive system load on remote integrations when executing actions. A bigger batch size might create overhead on the integration thus failing all entries. So, you need to be careful when increasing this value.

Max Postpone

Specifies the maximum number of action retries. If an action cannot be executed for any reason, such as connection failures, authentication problems or another SOAR internal problem, it will automatically be retried later. There are a number of global configuration parameters to configure how and when it will retry, but, after a number of retries specified in this field, SOAR will give up and mark the action as failed. Default value is 6 (in hours).

Connection Limit

Specifies the maximum number of concurrent connections for the integration. Default value is 5.

Max Action Retry

Specifies the maximum action retry count for the integration. Default value is 5

Max Rollback Retry

Specifies the maximum rollback retry count for the integration. Default value is 5.

Editing and Deleting an Integration

You can edit an existing integration by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Integration Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

You can delete an existing integration by clicking the **Delete** button under the **Actions** column.

For detailed information see the related Integration Guides.

Testing Integration

Click **Test** to verify the integration configuration.

When you click the **Test** button, it triggers the availability check for integration and if anything fails, a detailed error message is displayed. For example, in the case of a Check Point Firewall integration, SOAR needs a credential to work with the integration. If a credential is not available, an error message is displayed.

If the administrator of the remote integration accidentally deletes the credential that SOAR uses, SOAR will no longer be able to create actions on the integration. In this case, the integration is shown as offline (and an internal alert is created) and the error message is logged into the error log.

You can click **Test** button to see the error message.

A successful test marks the integration as **online**.

Flushing Queues

To flush the ques, select **Flush Queue** button under the **Actions** column of the integrations list. Following is the basic flow in SOAR:

1. Alert is received.

2. Matched playbooks run.

3. Action and rollback queue objects are created (waiting for execution in the queue).

4. Actions/rollbacks in the action/rollback queues are executed and saved.

When you click the **Flush Queue** button, SOAR starts executing actions/rollbacks without waiting for the execution scheduler (which consumes action/rollback queue objects).

Configuring Additional Parameters

Select **RESPOND > Configuration > Parameters**.

You might require performing some additional configuration, depending on the component or platform integration requirements. The **Parameters** tab displays a list of parameter that can be used for the additional configuration. For more information about additional configuration for integrations, see the respective Integration Guides.

- Searching a Parameter
- Editing a Parameter

Searching a Parameter

You can search a specific parameter, through the **Search** option. Click the button next to search, to view search results based on the following attributes:

- Parameter Name
- Parameter Value
- Default Value
- Description
- Last Modified By
- Modification Date
- Actions

Editing a Parameter

You can edit an existing parameter by clicking the **Edit** button under the **Actions** column. In the **Configuration Editor** window, specify the details for the following fields:

Parameter: Specify the parameter name.

Value: Specify the parameter value.

Description: Specify the parameter description.

Default Value: Specify the default value of the parameter.

You can not delete a parameter as it can be used in several integrations.

Configuring Case States

Select **RESPOND > Configuration > Cases**.

Fusion in the ArcSight Platform User's Guide enables you to customize case states such as statuses, severities, types and labels as per your requirement. You can configure multiple options to define case states to suit your requirement.

When you click the **Cases** page, you can view the following sub tabs:

- Statuses
- Severities
- Types
- Labels

Configuring Users

Select **RESPOND > Configuration > Users**.

An admin can list and edit user roles who will be interacting with ArcSight SOAR for case handling. SOAR authenticates users from Platform's single sign on provider. Initially, on authentication all the users are assigned with a **Super user** role, which can be later modified to the respective suitable role by the admin.

The **Users** page displays the list of users with options to modify an existing one.

- Searching for a User
- Editing a User

Searching for a User

You can search a specific user, through the **Search** option. Click the ![gear button] button next to search, to view search results based on user's **ID, User Name, Last Modified By, Modification Date, External User, Active User, User Role** and **Actions** filters.

Editing a User

You can modify an existing user's role, phone number and avatar by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **User Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

Configuring User Groups

Select **RESPOND > Configuration > User Groups**.

The **User Group** page displays a list of user groups and provides you options to create more user group or edit an existing user group.

- Searching for a User Group
- Creating a User Group
- Editing and Deleting a User Group

Searching for a User Group

You can search for a specific user group, through the **Search** option. Click the button next to search, to view search results based on list's **Name, Content Type, Size, Action Allowed, Enrichment Allowed, Last Modified By, Modification Date** and **Actions** filters.

Creating a User Group

Click the **+ Create User Group** button to create a new user group. In the **User Group Editor** window, specify the details for following fields:

Value	Description
Name	Name of the user group. Provide an explanatory name which gives a hint about what the user group is created for.
Email	Specify an email ID for the user group.
Users	Select the users to be included in this user group.
Avatar	You can select an avatar for the group by clicking on the **Choose File** button. Any image will work. It is recommended to select image files with sizes of 200 x 200 pixels.

Editing and Deleting a User Group

You can edit an existing user group by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **User Group Editor** window is displayed. Specify the details as per your requirement and click **Save**.

You can remove an existing list by clicking the **Delete** button under the **Actions** column.

Configuring Roles

Select **RESPOND > Configuration > Roles**.

The **User Role** page displays a list of user roles and options to create and edit them. The user roles define the permissions granted to a user. Your role determines which features that you can access. Following are the predefined roles:

Default Roles	Permissions
Admin	All Playbook permissions All Dashboards and Reporting permissions All Status permissions All Configuration permissions
Integration Owner	View Dashboards View Logs Manage Integration Configurations Manage Integration Credentials
Analyst	All cases permission All Playbooks permission All Dashboards and Reporting permission except for Manage Report Templates Status- View Alerts, Manage Actions and Rollback Queue, Manage Alerts and View Action and Rollback Queue permissions Configurations: From Alert Sources and Integrations-View Alert Sources, View Integration Configuration, and View Integration Credentials permissions From Security- View Users and View User Groups permission From Lists- View Exclusion Lookup Tables and View Lookup Table permissions
Super User roles	All permissions

A new user is automatically assigned the role of a **Super User** on the first log in. The admin can then choose to reassign a new role to the newly authenticated user.

A user cannot be assigned more than one role.

Creating a User Role

Click the **Create User Role** button to create a new user role. In the **Role Editor** window, specify the user role attributes as follows:

Value	Description
Role Name	Name of the user role. Consider giving an explanatory name that hints about the permission level of the user, such as., Full Administrator or Monitoring Operator.

Editing and Deleting a User Role

You can edit an existing user role by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Role Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

> You can not modify a **Super User** role.

You can also remove an existing user role by clicking the **Delete** button under the **Actions** column.

Configuring Access Control Lists

Select **RESPOND > Configuration > Access Control Lists**.

The Fusion in the ArcSight Platform User's Guide provides **Access Control Lists** (ACLs) to control the access of the users or user groups on SOAR objects. These objects include action capabilities, credentials, custom scripts, enrichment capabilities, enrichment plugins, integrations and integration types. For example, you might prefer a specific group of users to access some specific integrations. In such scenarios you can edit the access controls of the user groups in **Access Control Lists** tab.

When you click **Access Control Lists** tab, a list of SOAR objects along with users or user groups who can access those objects, and the last user and last modification date of an access control is displayed.

- "Searching an Access Control List" below
- Editing and Resetting an Access Control List

Searching an Access Control List

You can search a specific access control list, through the **Search** option. The search list keeps getting updated as you type. Click the ⚙▾ button next to **Search**, to view search results based on **Object, Access, Last Modified By** filters.

Editing and Resetting an Access Control List

You can edit an existing access control list by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Access Control List Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

You can only edit the **Allow Access For** and **Users/Groups** fields. By default, the object list is created with the **Anyone** option. If you want to narrow down the users for an object, just edit the related object and specify the users or groups for the option you selected in the **Allow Access For** field, which can be **Only selected users/groups** or **Anyone except selected users/groups**.

After editing an **Access Control List** item, **Clear** button appears in the **Actions** column of that item. You can not remove an **Access Control List** item from the list. By clicking on the **Clear** button under the **Actions** column, you can reset value of the **Access** column.

Setting Up SOAR for Customization

Select **RESPOND > Configuration > Customization Library**.

You can customize case management through plugin scripts, email templates, query templates, scriptable integration codes, text and HTML templates (used for notifications and other capabilities) and other customizations. The **Customization Library** tab displays a list of customzation, and also allows you to create a new customization content or modify an existing one. When a new plugin is uploaded through configuration/integrations/upload plugin options, you can also view and manage its code on this tab.

The list of customization can be filtered to display all integrations customizations, all integration types customization and all script types customization.

- "Searching a Customization" below
- "Creating a Customization" below
- "Filtering Integration Customizations" on the next page
- "Filtering Integration Types Customizations" on the next page
- "Filtering Script Types Customizations" on the next page
- "Editing, Deleting and Reseting a Customization" on the next page

Searching a Customization

You can search a specific user customization, through the **Search** option. Click the [gear icon] button next to **Search**, to view search results based on **ID, Name, Script Type, Integration Types, Integration, Last Modified By, Modification Date** and **Actions** filters.

Creating a Customization

Click the **+Create New Customization** button to create a new customization. In the **Customization Editor** window, specify customization name, description and type and enter the respective code in the black console.

Filtering Integration Customizations

Click **Show all integrations** filter to view the customization list based on the visible name of the integrations already defined on the environment.

Filtering Integration Types Customizations

The **Show all integration types** filter enables you to filter the list based on integration/plugin type.

Filtering Script Types Customizations

When you click **Show all scripts type** filter, a list of script type customizations is displayed.

Editing, Deleting and Reseting a Customization

You can edit an existing customization configuration by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Customization Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

You can delete an existing customization configuration by clicking the **Delete** button under the **Actions** column.

The **Reset** button resets the content of customization to out-of-the-box version.

The **Lookup** button shows the details about where this particular customization is being used.

Referencing Documents

Select **RESPOND > Configuration > Document Repository**.

The system can store your reference documents that might be linked to cases if needed. For example, you can add case handling guides for your SOC analysts and link these documents automatically when a case is created on SOAR.

- "Searching a Document" below
- "Uploading a Document" below
- "Editing and Deleting a Document From The Repository" below

Searching a Document

You can search a specific document, through the **Search** option. Click the button next to search, to view search results based on document's **ID, File Name, Title, Description, File Size**and **Actions** filters.

Uploading a Document

Click the **+Upload Document** button to upload a new document in the repository. In the **Document Repository Editor** window, specify details such as document **Title**, **Description** and then select the file to be uploaded.

Editing and Deleting a Document From The Repository

You can edit an existing document by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Document Repository Editor** window is displayed. Specify the details as per your requirement and click **Save**.

You can delete an existing document by clicking the **Delete** button under the **Actions** column.

To download a document click the **Download** button under the **Actions** column.

Setting Up Scope Items

Select **RESPOND > Configuration > Scope Item Property**.

With ArcSight SOAR 3.4, it is possible to create your own scope item property types and use them in your workflows. In order to define a new scope item property, click **Create Scope Item Property** button and specify the **Name** and **Data Type** fields. Scope item property can possibly be a:

- Number
- Text
- Json
- Percentage
- Boolean

Understanding Cases User Interface

SOAR has a very user friendly interface for tracking, viewing and managing cases in a single pane of glass. The **Case** tab enables you to perform multiple operation on one page. You can view the list of cases, edit the case information such as its label, status and priority, add and edit assignees, add watchers , related cases, comments, and attach files, all on a single **Case** page. To perform a deeper analysis, you can also fetch enrichment for cases and perform desired actions. The Case Management service desk also facilitates the flexibility to create manual cases and generate reports for analysis.

A typical **Case** page is displayed as follows:

Filtering Alerts For Case Creation

Select **RESPOND > Playbook > Rule Name Filters**.

At this stage, no case is formed. The rule name filters are used to decide the plan of action for an alert. The rules in this tab decide whether to register a case with the alert or not.

For example, if an alert is a possible threat, you can create a case with the alert, or you can receive the alert, save it and create a case but ignore all base events, or you can completely ignore the alert.

The **Rule Names** are listed in the ascending order in the **Rule Name Filters** page.

- Creating an Alert Source Rule Name Filter
- Managing Scope Item Extraction
- Searching for a Rule Name
- Editing an Alert Source Rule Name Filter

Creating an Alert Source Rule Name Filter

To create an Alert Source Rule Name Filter, complete the following procedure:

1. Click the **Create Alert Source Rule Name** button.
2. Specify a value for Rule Name. Make sure that the specified rule name matches with the correlation event name.
3. In the **Alert Source** menu, select an alert source from the list of created alert sources.

4. In **Ignore Mode** menu, select one of the following values:

Parameter	Description
Create alerts	When an alert with this rule name is received, a case is created irrespective of the alert sources and alert source types. All base events will be fetched if possible.
Ignore base events	It creates a case, but ignores base events if there are any
Ignore for all alert sources	It does not create cases for this rule name, irrespective of alert sources defined on the system.
Ignore for all alert sources of this type	It does not create a case when an alert with this rule name is received, only for the alert sources of the type shown in the **Alert Source Type** field. It creates cases for the alert sources of other types.
Ignore for this alert source	It does not create a case when an alert with this rule name is received, only for the alert source shown in the **Alert Source** field. It creates cases for the other alert sources.

5. Click **Save**.

Managing Scope Item Extraction

You cannot edit an existing *Scope Item Extraction*. You can delete an existing *Scope Item Extraction*. To delete, click **Delete** in the **Actions** column. The Scope Item Extraction Section has the following information:

Parameter	Description
Field Name	Name of the field.
Select Source	Select from the list [**Base Event, Correlated**].
Select Category	Select from the list [**Computer Name, Email Address, File, File Name, Hash, Host, Keyword, MAC Address, Network Address, Process, Rule name, Unknown, URL, Username**].
Select A Role	Select from the list [**Impact, Offender, Related**].
Add	Click this button to add the scope item.

Searching a Rule Name

You can search a specific Rule Name. Click the button next to search, which allows you to view search results based on **Rule Name ID, Rule Name, Alert Source, Ignore Mode, Pattern Matcher**, and **Actions**.

Editing an Existing Alert Source Rule Name Filter

You can modify an existing Alert Source Rule Name Filter to configure additional extraction from the base events or the correlated events.

You cannot rename or delete an existing alert source rule name filter.

To edit an existing Alert Source Rule Name Filter, complete the following procedure:

1. Do one of the following:

   - Double-click the record that you want to edit.

   - Click **Edit** in the **Actions** column of the related record.

   For example, to edit the Alert Source Rule Name Filter for Real-time Threat Detection, make sure that the Alert Source name is Real-time Threat Detection.

2. In **Ignore Mode** menu, select a value according to your needs.

   You will be able to see the created tickets in the **Cases** page.

Classifying Cases on SOAR

Select **RESPOND > Playbooks > Classification**.

Classification tab helps you to organize or maintain your cases on the SOAR platform.

After an alert is received and a case is created with it, then it is passed for classification. On **Classification** tab, the alert is labeled depending upon the conditions. The rule names are checked depending on the rule name and a label is added to the alert. Later these classification labels help the system in choosing and executing a playbook for this case.

You can view a list of classification on this page. The Classification list is processed from top to bottom and only the first match is executed. You can edit the rank of a classification rule through the **Rank** option and created items will appear as the last item in the table.

- "Searching a Classification" below
- "Creating a Classification Rule" below
- "Editing and Deleting a Classification" on page 275

Searching a Classification

You can search a specific **Classification** through the **Search** option. Click the ⚙️▾ button next to search to view search results based on **Classification ID, Rule Conditions, Rule Actions, Last Modified by, Modification Date, Rank** and **Actions**.

Creating a Classification Rule

You can create a classification with no condition, which will execute on all cases. You cannot create a classification without any action. After you select a condition, SOAR matches it with the alert conditions and automatically creates actions, that is defined under the **Actions** field.

Click the **Create Classification Rule** button to create a new classification. In the **Classification Editor** window, specify the details for following fields:

Matching Mode: Select <All condition, Any condition> to specify if the new rule allows all or any condition to be matched , similar to a logical AND /OR mode.

Create Conditions:

- **Type**: Select a condition type from the drop-down list. Following table presents the detailed condition types:

Type	Description
Address contains	An address value which will be searched in the IP address of alert sources. You can use the "*" character as the wildcard. Assume that this value is .*.*.22, then the condition will be met when a case is created for all the alert sources having IP addresses that end with "22".
Address doesn't contain	Condition will be met when the value typed here is not a part of alert source IP addresses.
Address is in subnet	A subnet value, which will be searched in the subnet address of alert sources. You can use the __*__ character as the wildcard.
Address is not in subnet	Condition will be met when the value typed here is not a part of alert source subnet addresses.
Address matches regex	Condition will be met when the IP address of the alert source is matched to the regular expression specified here.
Address doesn't match regex	Condition will be met when the IP address of the alert source does not match the regular expression specified here.
Alert is manual	Condition will be met when the alert is created manually.
Alert is not manual	Condition will be met when the alert is not created manually.
Alert parameter matches key value pair	Pair can be given as key=value. Condition will be met when the parameter (key) is equal to the value specified here for any alert parameters.
Alert parameter doesn't match key value pair	Condition will be met when the parameter (key) is not equal to the value specified here for any alert parameters.
Alert source is	Condition will be met when the alert source of the related case is the one selected here.
Alert source is not	Condition will be met when the alert source of the related case is not the one selected here.
Alert source rule name is any of	Condition will be met when the rule name of case's alert source is any of the selected values here. You can select multiple rule names in the **Parameters** combo box.
Alert source rule name is not any of	Condition will be met when the rule name of case's alert source is not any of the selected values here. You can select multiple rule names in the **Parameters** combo box.

Type	Description
Alert source rule name is in list	Condition will be met when the alert source rule name of the related case is in the list selected here.
Alert source rule name is not in list	Condition will be met when the alert source rule name of the related case is not in the list selected here.
Alert source rule name matches regex	Condition will be met when the alert source rule name is matched to the regular expression specified here.
Alert source rule name doesn't match regex	Condition will be met when the alert source rule name is not matched to the regular expression specified here.
Alert time is between (day of week)	Condition will be met when the creation time of an alert is between the dates and times selected here.
Alert time is not between (day of week)	Condition will be met when the creation time of an alert is not between the dates and times selected here.
Alert time is between (time of day)	Condition will be met when the creation time of an alert is between the times of each day selected here.
Alert time is not between (time of day)	Condition will be met when the creation time of an alert is not between the times of each day selected here.
Assignee is	Condition will be met when the assignee of the related case is the one selected here.
Assignee is not	Condition will be met when the assignee of the related case is not the one selected here.
Assignee is set	Condition will be met when the assignee of the related case is set.
Assignee is not set	Condition will be met when the assignee of the related case is not set.
Assignee is a member of group	Condition will be met when the assignee of the related case is a member of the group selected here.
Assignee is not a member of group	Condition will be met when the assignee of the related case is not a member of the group selected here.
Classification contains	Condition will be met when the classification typed here is in classification list.

Type	Description
Classification doesn't contain	Condition will be met when the classification typed here is not in classification list.
Scope item category is	Condition will be met when the scope item category of the related case is the one selected here.
Scope item category is not	Condition will be met when the scope item category of the related case is not the one selected here.
Scope item role is	Condition will be met when the scope item role of the related case is the one selected here.
Scope item role is not	Condition will be met when the scope item role of the related case is the one selected here.
Scope item value equals	Condition will be met when the scope item value of the related case is equal to the value expressed here.
Scope item value doesn't equal	Condition will be met when the scope item value of the related case is not equal to the value expressed here.
Scope item value is in list	Condition will be met when the scope item value of the related case is in the list selected here.
Scope item value is not in list	Condition will be met when the scope item value of the related case is not in the list selected here.
Severity is	Condition will be met when the severity of the related case is the one selected here.
Severity is not	Condition will be met when the severity of the related case is not the one selected here.
Status is	Condition will be met when the status of the related case is the one selected here.
Status is not:	Condition will be met when the status of the related case is not the one selected here.

- **Parameters**: Appropriate value for the type. Select from the list or enter a value.

Create Actions:

- **Action**: Select an action from **Add case label** and **Change severity of Case**.
- **Parameters**: Appropriate value for the type. Select from the list or enter a value.

> The newly created Classification Rule is displayed on the Classification page and is in **Disabled** state by default. You must ensure enabling the rule before using it.

Editing and Deleting a Classification

You can edit an existing classification by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Classification Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

You can delete an existing classification by clicking the **Delete** button under the **Actions** column.

> You cannot edit an existing condition or action. You have to delete the condition or action and create a new one.

Consolidating Alerts to Create Cases

Select **RESPOND> Playbook> Consolidation**.

Multiple alerts are generated from different alert sources that are integrated with SOAR. These alerts are automatically consolidated to create a case as per the configuration settings. The **Consolidation** page displays a list of rules to consolidate alerts to create cases.

When an alert reaches the consolidation plugin based on the rules, all the correlated alerts are consolidated to create a case. It is after this consolidation procedure that the system decides whether to create a new case or to add the alert into an existing one.

Consolidation rules are processed from top to bottom and only the first match is executed. Any alerts that matches the same consolidation rule is gathered in to the same case until that case status is **Close**. In that instance, a new case will be created and alerts are consolidated into this case.

- "Searching a Consolidation Filter" below
- "Creating a Consolidation Filter" below
- "Editing and Deleting a Consolidation Filter" on the next page

Searching a Consolidation Filter

You can search a specific **Consolidation Filter**, through the **Search** option. Click the ⚙▾ button next to search, to view search results based on **ID, Rule Conditions, Timespan, Last Modified by, Modification Date, Rank** and **Actions**.

Creating a Consolidation Filter

Click **Create Consolidation Filter** to create a new consolidation filter. In Consolidation Filter , specify the details for following fields:

Timespan: Value in minutes, hours, weeks or days. Timespan provides time intervals to consolidate alerts into one case.

Since Last Alert: Timespan will be calculated from the last alerts creation time.

Since First Alert: Timespan will be calculated from the first alerts creation time.

Until First Response: Consolidation will stop when the case is responded by an analyst. When this checkbox is selected, Fusion in the ArcSight Platform User's Guide will track the response status of the case and timespan and stop the consolidation at whichever comes first.

Create Conditions: Select a condition for alert consolidation from the following list of condition **Types** and **Parameters**:

- **Type**: Type of the consolidation. Select from the list.
  - Alert source is
  - Alert source rule name is any of
  - Alert source rule name is in list
  - Alert source rule name matches regex
  - Scope item category is
  - Scope item role is
  - Scope item value does not equal
  - Scope item value equals
  - Scope item value is in list
  - Scope item value is not in list
- **Parameters**: It varies depending on selected consolidation type.

> The newly created Consolidation Filter is displayed on the Consolidation page and is in **Disabled** state by default. You must ensure enabling the **Consolidation Filter** before using it.

Editing and Deleting a Consolidation Filter

You can edit an existing consolidation filter by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Consolidation Filter** window is displayed. Specify the values as per your requirement and click **Save**.

You can delete an existing consolidation filter by clicking the **Delete** button under the **Actions** column.

Dispatching Cases

Select **RESPOND > Playbook > Dispatch**.

You can define a set of dispatch rules to automatically assign a case to a user or user role or a tier. After consolidation, once a case is created, you can decide to assign it to a group/ a team or a person and also add a severity to the case. If you did not assign the case to any user/group/team, the system automatically selects the playbook, based on rules and labels, and then executes it to resolve the issue.

The Dispatch page presents a list of dispatch rules that must be executed for the cases with specified conditions.

Dispatch rules are processed from top to bottom and only the first match is executed. You can view the rank of the rule (to see the order of dispatch actions to be applied to the cases), conditions of the dispatch rule, dispatch actions, and the user and date of last edits performed on the rule.

You can edit the rank of a dispatch rule through the **Rank** column and created items appears as the last item in the table.

If you do not want to remove the rule permanently, you can disable it using the **Disable** button in the list.

- "Searching a Dispatch Rule" below
- "Creating a Dispatch Rule" below
- "Editing and Deleting a Dispatch Rule" on page 282

Searching a Dispatch Rule

You can search a specific **Dispatch Rule**, through the **Search** option. Click the ⚙▾ button next to search, to view search results based on **ID, Rule Conditions, Rule Actions, Last Modified by, Modification Date, Rank** and **Actions**.

Creating a Dispatch Rule

You can create a dispatch rule with no condition, which will execute on all cases. You cannot create a classification without any action. Once you select a condition, SOAR matches it with the alert conditions and automatically creates actions, that is defined under the **Actions** field.

Click **Create Dispatch Rule** button to create a new dispatch rule. In the **Dispatch Editor** window, specify the details for following fields:

Matching Mode: Select <All condition, Any condition> to specify if the new rule allows all/any the conditions to be matched , similar to a logical AND /OR mode.

Create Conditions: : To create conditions for the rule, click on the **Create** button within the **Conditions** box.

- **Type**

  Select the condition type from the **Type** drop-down list. Following table presents the detail condition types:

  Table: Condition Types

Type	Description
Address contains	An address value which will be searched in the IP address of alert sources. You can use the "*" character as the wildcard. Assume that this value is .*.*.22, then the condition will be met when a case is created for all the alert sources having IP addresses that end with "22".
Address doesn't contain	Condition will be met when the value typed here is not a part of alert source IP addresses.
Address is in subnet	A subnet value, which will be searched in the subnet address of alert sources. You can use the __*__ character as the wildcard.
Address is not in subnet	Condition will be met when the value typed here is not a part of alert source subnet addresses.
Address matches regex	Condition will be met when the IP address of the alert source is matched to the regular expression specified here.
Address doesn't match regex	Condition will be met when the IP address of the alert source does not match the regular expression specified here.
Alert is manual	Condition will be met when the alert is created manually.
Alert is not manual	Condition will be met when the alert is not created manually.
Alert parameter matches key value pair	Pair can be given as key=value. Condition will be met when the parameter (key) is equal to the value specified here for any alert parameters.
Alert parameter doesn't match key value pair	Condition will be met when the parameter (key) is not equal to the value specified here for any alert parameters.
Alert source is	Condition will be met when the alert source of the related case is the one selected here.

Type	Description
Alert source is not	Condition will be met when the alert source of the related case is not the one selected here.
Alert source rule name is any of	Condition will be met when the rule name of case's alert source is any of the selected values here. You can select multiple rule names in the **Parameters** combo box.
Alert source rule name is not any of	Condition will be met when the rule name of case's alert source is not any of the selected values here. You can select multiple rule names in the **Parameters** combo box.
Alert source rule name is in list	Condition will be met when the alert source rule name of the related case is in the list selected here.
Alert source rule name is not in list	Condition will be met when the alert source rule name of the related case is not in the list selected here.
Alert source rule name matches regex	Condition will be met when the alert source rule name is matched to the regular expression specified here.
Alert source rule name doesn't match regex	Condition will be met when the alert source rule name is not matched to the regular expression specified here.
Alert time is between (day of week)	Condition will be met when the creation time of an alert is between the dates and times selected here.
Alert time is not between (day of week)	Condition will be met when the creation time of an alert is not between the dates and times selected here.
Alert time is between (time of day)	Condition will be met when the creation time of an alert is between the times of each day selected here.
Alert time is not between (time of day)	Condition will be met when the creation time of an alert is not between the times of each day selected here.
Assignee is	Condition will be met when the assignee of the related case is the one selected here.
Assignee is not	Condition will be met when the assignee of the related case is not the one selected here.
Assignee is set	Condition will be met when the assignee of the related case is set.
Assignee is not set	Condition will be met when the assignee of the related case is not set.

Type	Description
Assignee is a member of group	Condition will be met when the assignee of the related case is a member of the group selected here.
Assignee is not a member of group	Condition will be met when the assignee of the related case is not a member of the group selected here.
Classification contains	Condition will be met when the classification typed here is in classification list.
Classification doesn't contain	Condition will be met when the classification typed here is not in classification list.
Scope item category is	Condition will be met when the scope item category of the related case is the one selected here.
Scope item category is not	Condition will be met when the scope item category of the related case is not the one selected here.
Scope item role is	Condition will be met when the scope item role of the related case is the one selected here.
Scope item role is not	Condition will be met when the scope item role of the related case is the one selected here.
Scope item value equals	Condition will be met when the scope item value of the related case is equal to the value expressed here.
Scope item value doesn't equal	Condition will be met when the scope item value of the related case is not equal to the value expressed here.
Scope item value is in list	Condition will be met when the scope item value of the related case is in the list selected here.
Scope item value is not in list	Condition will be met when the scope item value of the related case is not in the list selected here.
Severity is	Condition will be met when the severity of the related case is the one selected here.
Severity is not	Condition will be met when the severity of the related case is not the one selected here.
Status is	Condition will be met when the status of the related case is the one selected here.
Status is not:	Condition will be met when the status of the related case is not the one selected here.

- **Parameters**: Appropriate value for the selected condition type. Select from the list or enter a value.

Create Actions:

- **Action**: Defines case dispatch actions for the rule. Select the action from the **Action** combo box. Following are the available actions:

  - **Add a case label**: When selected, **Parameters** field toggles to a combo box listing the case labels defined in the system. You can choose a label from the list, so that when the case meeting the above conditions is created, it will be labeled as the one selected here.

  - **Assign to a user or group**: When selected, **Parameters** field toggles to a combo box listing the users/groups defined in the system. You can choose a user or group from the list, so that when the case meeting the above conditions is created, it will be assigned to the user or group selected here.

  - **Change severity of case**: When selected, **Parameters** field toggles to a combo box listing the case severities defined in the system. You can choose a severity from the list, so that when the case meeting the above conditions is created, the cases initial severity will be changed to the one selected here.

    Click the **Save** button within the **Actions** box to add your rule action. You can add as many actions as you want.

    > You cannot edit a previously created conditions or actions. You have to delete and create a new condition and action.

- **Parameters**: Appropriate value for the selected action type. Select from the list or enter a value.

> The newly created Classification Rule is displayed on the Classification page and is in **Disabled** state by default. Enable the rule before using it.

Editing and Deleting a Dispatch Rule

You can edit an existing dispatch rule by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Dispatch Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save**.

You can delete an existing dispatch role by clicking the **Delete** button under the **Actions** column.

> You cannot edit an existing condition or action. You have to delete the condition or action and create a new one.

Working with Playbooks

Select **RESPOND > Playbooks > Playbooks**.

A **playbook** defines the automation and orchestration capability. After a case is dispatched, playbook performs the response procedure. The system can execute a fully automated playbook as well as a semi-automated playbook.

A completely automated playbook does not require any decision making from the agents. A semi-automatic model requires agent intervention for decision making or providing some extra information to the automation. So during a semi-automation procedure, SOAR handles the case resolution automatically till some point and then the control is passed to agents for decision making task and again after the decision is made, the control is handled by automation. If needed, SOAR automation can again assign the task to agent for some decision making or extra information requirement. So basically, SOAR performs orchestration and then finally makes a **Response**.

You can specify the execution priority of playbooks by setting the **Rank** values for each playbook, the smaller the rank, the higher is the priority.

Playbooks are processed from top to bottom and when a case matches, all of the playbooks with matching conditions are executed.

While designing any playbook, you must set conditions to ensure if multiple playbooks can run on the same case or not. As the playbooks running on the same case are not aware of each other, they must be designed independently such that one playbook does not interfere with another. If possible, it is recommended that a case matches with only one playbook.

- "Searching a Playbook" on the next page
- "Creating an Advanced Playbook" on the next page
- " Creating Workflow Playbook" on page 285
- "Executing Workflow Playbooks" on page 285
- "Workflow Playbook Elements" on page 285
- "Types of Connectors in the Workflow Playbook" on page 289
- "Importing and Exporting a Workflow" on page 289
- "Editing Rank of a Playbook" on page 289
- "Editing and Deleting a Playbook" on page 289

Searching a Playbook

You can search a specific **Playbook**, through the **Search** option. Click the ⚙▾ button next to search, to view search results based on the following attributes:

- ID
- Scenario Name
- Type
- Last Modified by
- Modification Date
- Rank
- Actions
- Disabled

Creating an Advanced Playbook

The **Advanced Playbook** allows you to write your own playbook scripts.

1. Click **Create Advanced Playbook** button.
2. In the **Advanced Playbook Editor** window, specify the details for following fields:

Value	Description
Name	Display name of the playbook.
Matching Mode	**All Conditions** means playbook will be executed if all the conditions are true. **Any Conditions** means playbook will be executed if any of the conditions is true.
Rollback Mode	Set if the action will be permanent or will be rolled back after a period of time.
case auto-close	From the combo box, you can select in which conditions the playbook will close the cases.
Conditions	Click **Create** to add a condition to this playbook. You can define multiple conditions.

3. In the black console area, you can write your playbook scripts in Python programming language.

4. To test your playbook, use the **Test** option:

   a. Select a defined alert source from the combo box.

   b. For **Value to Block**, enter a value to test your script.

   > The option **Value to Block** can be any parameter depending on your script, such as IP or email address.

   c. Click **Test**.

   Your test result is displayed on the same console.

Creating Workflow Playbook

Workflow Playbooks run automatically and follows the visual process definition. You can specify the a name to the playbook in **Playbook Name**.

While creating a Workflow Playbook, you can drag and drop elements from the right side of the page. You must enter appropriate and valid values depending on the element in the **Properties** tab. Each element must be connected to another except the last one.

When a case is created, a playbook with matching condition is executed. The match conditions of the Workflow Playbook are defined in the **Start** element of the playbook.

Executing Workflow Playbooks

Workflow Playbooks are run automatically when:

- **A new case is created**: cases are created by the Alert Rule Name Filter configuration.
- **A new alert is received**: Alerts are added to the cases by the Consolidation rules.
- **Rules of the case is updated**: Some alert sources update an existing alert for example, QRadar Offences and these can trigger an execution.

Workflow Playbook Elements

To create a visual process definition, you must map the executable instructions through the predefined workflow playbook elements. You can drag and drop following elements to create the workflow:

- **Automation Bit Usage**: Automation Bits are custom code created by the users to execute custom business logic. A detailed explanation for Automation Bit's can be found in

Automation Bit section of this guide. While using bits, scope will be supplied from the **Start from here** element if **Scope Filter** variable is not used.

- **Actions Usage**: There are two kinds of actions:

  - Actions coming from the SOAR itself, and these actions act on cases to change it appropriately, for example, Status, Severity.

  - The action capabilities coming from integrations. There are different capabilities depending on the target device and all of them takes some input regarding their role in the workflow.

    Action elements are named as <Integration Name> - <Capability Name>. For example, Active Directory - Lock User.

  Actions usage have several standard properties including:

Properties	Descriptions
Title	Visible name of the element in the visual editor.
Continue on Error	In some cases an action on a device can return an error for example, network problems. In such cases , SOAR will stop the execution of the workflow entirely. If this option is selected, SOAR will continue execution even if the action has failed.
Rollback Mode	SOAR can undo the action after a set time if needed. In many devices there are limits to how many items can be blocked and most of these artifacts usefulness drops over time. Rollback future gives the SOAR users a way to control their actions and the health of the target device.
Scope Filter	The scope filter name can be changed from capability to capability but in essence filter will define which scope items from the alert will be included in the execution. Some actions also have other fields and these are populated from data that resides on the target device. Such as tag's or group names.
Actions are synchronous	Therefore when a workflow processes an action element, it queues this action and after successful queueing of this action workflow will resume processing the next element. This means in an ideal SOAR, processing actions will not create a performance issue for the workflow execution. There however some edge cases that when SOAR is under heavy load or an unexpected error is present, actions might be queued but different elements are executed before these actions are finished.

- **Enrichment**: Enrichments are data gathering capabilities that will assist in case response procedures and decision making.

  Enrichments have several standard properties including:

Properties	Descriptions
Title	Visible name of the element in the visual editor.
Continue on Error	In some cases an enrichment on a device can return an error e.g network problems. In such cases SOAR will stop the execution of the workflow entirely. If this option is selected SOAR will continue execution even if the enrichment is failed.

Properties	Descriptions
Integration	On which integration this capability will be executed.
Scope Filter	This part's name can be changed from capability to capability but in essence filter will define which scope items from the alert will be included in the execution.
Do not use cache	When a workflow processes an action element, it queues this action. After successful queueing workflow resumes processing the next element. This means in an ideal SOAR, processing actions does not create a performance issue for the workflow execution. However some when SOAR is under heavy load or an unexpected error is present, actions might be queued but different elements are executed before these actions are finished.
Enrichments are synchronous	When executed they will start immediately and hold the workflow execution on this state until a result is returned. It is important to note that not every enrichment works as fast as you expect and in some cases rate limits might apply affecting the execution time of the overall workflow. Some enrichments execute and then wait for the process to be completed in the target device. These are also called asynchronous for their update part but for workflow execution they are treated as synchronous as well and will stop the execution until the response is returned.

- **Tasks**: Tasks are elements that does not have an automatic component. These elements are dependent on SOC analysts for completion. Task properties are dependant on the configuration of the task. So one or more of these properties might not appear in **Properties** tab..

  Tasks have several standard properties including:

Properties	Descriptions
Title	Visible name of the element in the visual editor.
Scope Filter	The name can be changed from task to task but in essence filter will define which scope items from the alert will be included in the execution. Filters can occur more than once and they are restricted to the Scope Item Type defined for them. So a **Network Address** type filter only works on **Network Address** type scope items.
Timeout Span	It is when the task is due, it will be defined by this property. Task will be timed out when it is due and execution will continue. If left empty, this value will be taken from the Configuration Parameter **WorkflowTimeout** as a global value.

- **Analyst's Decision**: This is the logic element and provides true/false options to the analyst.

  Analyst's decision have several standard properties including:

Properties	Descriptions
Title	Visible name of the element in the visual editor.
Description	Description of the decision.

Timeout span	This property is defined when the task will be due. When the task is due will be defined by this property. Task will be timed out when it is due and execution will continue. If left empty, this value will be taken from the Configuration Parameter **WorkflowTimeout** as a global value.
Send Additional Email for Approva	When this is checked, SOAR will send an additional email for out of SOAR interaction to the selected Analyst.
Analyst	Recipient of the approval Email.

- **Utilities**: There are three types of utility elements:

  ○ **Notification**:This element supports sending notifications to different users.

  Notifications can be sent from different channels and currently on-screen, SMS, email and windows type messages are supported. Notifications use free-form subject and a pre-defined template for the message.

  ○ **Decision**: Decision are standard logic element of the workflow. For a given predicate group in the property section, SOAR checks the alert scope and the workflow scope. If both of the scopes match, the automation returns a **true** value and the playbook is executed.

  The alert scope is defined at the **Start from here** element and workflow scope is the enrichment data that is specific to the workflow execution gathered till this point.

  ○ **User Decision**: User decisions are true/false type checkpoints and they are sent to a recipient for gathering inputs.

  The difference in the **Task Decision** and **User Decisions** can be explained as, the user decision sends the decision message to a variety of recipients. It can send the notification to a free-text e-mail address, to a user, or to an the case scope.

  User decision takes a template to form the message and expects the recipient to reply with an **APPROVE** or **DENY** option. You can create more than one template to send different set of data and messages to the relevant recipients. You can find the **User Decision Notification Email Template** as a built-in template in the **Customization Library**.

  You can also define scope restricted parameters that can be filled on the fly.

  > Using a scope restricted parameter in the e-mail subject shows only the first item in the parameter. Rest of the items are appended in the body of the message. The decision must appear in the body of the reply message.

Types of Connectors in the Workflow Playbook

Every element in workflow has a pre-defined connector type. There can be one, two or three output connectors.

- **Single connector**: All actions and most other types of elements, fall into this category and after the element executes workflow continue to the next element.
- **Double connector**: Elements that contain a timeout falls into this category. First connector will lead to a successful completion of the element within the given time, these are named **then** and second connector will lead to timeout.
- **Triple connector**: User and Analyst Decision falls into this category. First two connectors will lead to true and false respectively in a successful execution and third connector will lead to timeout.

Importing and Exporting a Workflow

You can import a pre-designed workflow by clicking the **Import Workflow** tab. In **Workflow Import Editor** window, navigate to the template file, add a suitable name for the template and then click **Save** to import a a workflow.

To export a workflow playbook, click **Export** option under the **Actions** tab.

> You can not export an advanced playbook

Editing Rank of a Playbook

You can define the order of execution for different playbooks by assigning a rank to it. Click **Edit Rank** option under the **Actions** tab and then modify the rank of the playbook in the respective **Rank** column.

Editing and Deleting a Playbook

To edit the previously created playbooks click **Edit** option under the **Actions** tab. In the **Workflow Playbook Editor** window, modify the visual process flow to suit you requirements.

To remove a playbook from the automation, click **Delete** option under the **Actions** tab.

Handling Repetitive Tasks With Scheduled Playbooks

Select **RESPOND > Playbooks > Scheduled Playbooks**.

You can use a **Scheduled Playbook** to close repetitive tasks or automate time-based mundane tasks.

- "Searching a Scheduled Playbook" below
- "Creating Scheduled Playbooks" below
- "Editing and Deleting a Scheduled Playbook" on the next page

Searching a Scheduled Playbook

You can search a **Scheduled Playbook**, through the **Search** option. Click the ⚙▾ button next to search, to view search results based on **ID, Name, Type, Description, Last Modified by, Modification Date** and **Actions**.

Creating Scheduled Playbooks

To create a new scheduled playbook, click the **Create Scheduled Playbook** button. In the **Scheduled Playbook Editor** window, specify the details for the following fields:

Value	Description
Name	Display name of the scheduled playbook.
Trigger Frequency	For Trigger Frequency, select from Every minute, Every 5 minutes, Every 10 minutes, Every 30 minutes, Every hour, Every 2 hours, Every 3 hours, Custom cron value (to define your own frequency) options.

In the console area, you can type a script for the playbook using Python programming language.

After typing the script, you can test the playbook using the **Test** option. Select a defined alert source from the combo box, type a value into the **Value to Block** field, and then click **Test** . Your test results are displayed on the same console.

> The option **Value to block** can be any parameter depending on your script, such as IP or email address.

You can also refer to the API Documents at the top right of the **Scheduled Playbook Editor** window.

Editing and Deleting a Scheduled Playbook

You can edit an existing scheduled playbook by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Scheduled Playbook Editor** window is displayed. Specify the values in editor window or edit the playbook script as per your requirement and click **Save** to modify.

You can delete a scheduled playbook by clicking the **Delete** button under the **Actions** column.

Creating Custom Business Logics

Select **RESPOND > Playbooks > Automation Bits**.

Automation Bits are custom code created that you create to execute custom business logic. ArcSight SOAR supports Python as programming language to write an automation bit.

- "Searching an Automation Bit" below
- "Creating an Automation Bit" below
- "Editing and Deleting an Automation Bit" on the next page

Searching an Automation Bit

You can search a specific **Automation Bit**, through the **Search** option. Click the ⚙️▾ button next to search, to view search results based on **ID, Name, Language, Last Modified by, Modification Date** and **Actions**.

Creating an Automation Bit

Click the **+Create Automation Bit** button to create a new automation bit. In the **Automation Bit Editor** window, specify the details for following fields:

Name: Name of the Automation Bit.

Description: Description of the Automation Bit.

Input Parameters: Starting parameters of the Automation Bit. These can be **Date**, **String** or **Scope Filter** and named here to be used in the Automation Bit. **Date** results in current time. **String** creates a parameter input field in workflow playbooks. **Scope Filter** creates a filter field in workflow playbooks.

Automation Bit's are **syncronous** and will hold the workflow executions until they are done.

> 🏠 This capability, if used in unexpected ways, might create longer than usual workflow execution times and delays.

You can type your **Automation Bit** script at the Black Console.

Editing and Deleting an Automation Bit

You can edit an existing automation bit by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Automation Bit Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

You can delete an existing automation bit by clicking the **Delete** button under the **Actions** column.

Managing Triggers

Select **RESPOND > Playbooks > Triggers**.

Triggers are mini playbooks that are triggered by several events. These events are created by human interaction or passage of time where SLA is concerned. Triggers evaluate the changes in the cases and if it matches to a trigger execution condition, the trigger starts automatically. Trigger executions are done from **top to bottom** and all triggers that matches the conditions will run. Only **Event Type** condition can be used in trigger **Start Condition** and the rest of the execution is done in the workflow through **Decision** elements.

As events can not be matched to two different **Event Type**, so **AND** operator is not supported.

- "Searching a Trigger" below
- "Creating a Trigger" below
- "Editing and Deleting a Trigger" below

Searching a Trigger

You can search a specific **Trigger**, through the **Search** option. Click the ⚙▾ button next to search, allows you to view search results based on **ID, Name, Last Modified by, Modification Date**, **Rank** and **Actions**.

Creating a Trigger

To create a trigger, click the **Create Trigger** button. In the **Trigger Playbook Editor** window, drag and drop the elements to create a workflow. To understand more on creating workflow, see **Creating Workflow Playbook.**

Editing and Deleting a Trigger

You can edit an existing Trigger by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Trigger Playbook Editor** window is displayed. Modify the properties of the Trigger Playbook elements or add or delete the element as per your requirement and then click **Save**.

Handling Manual Processes Through Tasks

Select **RESPOND > Playbooks > Tasks**.

Tasks are a way to define manual processes for Case response. The system can handle the automatic and manual elements together in a defined workflow. Analyst Task creates a task that is handled by the SOC Analysts within the SOAR Case Management.

- "Searching a Task" below
- "Creating a Task" below
- "Editing and Deleting a Task" on the next page

Searching a Task

You can search a specific **Task** through the **Search** option. Click the ⚙▾ button next to search, to view search results based on **Name, Description, Task Scopes, Task Output, Last Modified by, Modification Date** and **Actions**.

Creating a Task

You can define the **Analyst Tasks** in this window and the resulting task can then be used in the workflow as a standard element. To create a task, click the **+Create Analyst Task** button. In the **Analyst Task Editor** window, specify the details for the following fields:

Name: Visible name of the element in the visual editor.

Description: Description of the Task to be shown to the analyst.

Task Scope: Task scope is enabled here and these items will be filtered and shown to the analyst and expected to be completed by him/her.

Scope Item Categories: Input scope item types are selected here. This area supports multiselection.

Task Output: Task output is enabled here.

Scope Item Category: Expected scope item type is selected here. Scope item's created by the analyst will have this type. This area is single selection.

Task Merge: If in a case has more than one alert or a consolidation is ongoing it is possible that the workflow will run more than once and there will be tasks recurring for the analyst to complete. **Task Merge** gathers tasks occurring from the same workflow and shows them as one task to the analyst reducing their load. **Timeout Span** will be merged as well and SOAR will update the merged tasks **Due Time** as the most current one.

Using Task Output or Analyst Decision will disable **Task Merge** capability of SOAR for that elements. **Task Scope** is limited to handle **200** scope items. A task containing more than 200 scope items will be divided into more than one task.

Editing and Deleting a Task

You can edit an existing task by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **Analyst Task Editor** window is displayed. Specify the values in editor window as per your requirement and click **Save** to modify.

You can delete an existing task by clicking the **Delete** button under the **Actions** column.

Displaying Alerts

Select **RESPOND** > *Status* > *Alerts*.

To understand the SOAR system status, you can view all the alerts that the system has ingested within last 30 days. You can also customize the alerts display list using filter parameters.

Customizing the Alerts Display List

You can view selected alerts by selecting the appropriate filter option.

You can select an alert source in the **Alert Source** combo box and see the alerts only generated by the selected source. You can also narrow down the alert list by providing a time interval (Start/End Dates) and specific parameters (Alert Parameters) that are included in the alerts' context.

After selecting the alert filters, a list of alerts is displayed with following details:

Parameter Name	Description
ID	Alert ID
Created At	Date and time of the alert creation.
Alert Source	This is the visible name of the alert source. A visible name is assigned to an alert source during its configuration on SOAR platform.
Cases	The Cases related with the alert. When SOAR ingested this alert, this Case is created.
Selected Alert Parameters	Some of the parameters of the alerts.
Actions	Alert Details: Click the button to view alert details. In the **Alert Details** page you can view scope items associated with this alert along with the detailed information on alert source, time and date to create and update the alert.
	Show Parameter as Json: Click the button to view all the alert data in JSON format.
	Show Executed Playbooks: Click the button to view the playbook that was run for this alert.
	Process Again: Click the button to re-evaluate all the playbooks and if SOAR finds any new playbook or updated playbook with the matching condition, it will run the playbooks again for this alert.

Action and Rollback Queues

Select **RESPOND** > **Status** > **Action and Roll backQueue**.

SOAR has a mechanism to manage actions to be executed on the integration, called queuing. This section explains the action and rollback queues.

When SOAR receives an alert, alert is processed according to playbooks and SOAR decides the action and target integration.

SOAR adds this action process or rollback process to **Action and Rollback Queues** list which you can ignore approve or clear items. In order to filter list based on process type, Integration type, you can use buttons on the top of the list.

Action History

Select **RESPOND** > **Status** > **Action History**.

Action History tab lets you display and search logs of executed actions and rollback operations. To manage action history, click on the **Action History** tab in **Status** menu.

The page allows you to filter the action list by the following criteria:

- **Stage**: Stage of the action. Available values are **Executed Actions** and **Rollback Actions**.
- **Device**: You can select a device defined on your system to see the actions only performed on that device.
- **Playbook**: You can select a playbook defined on your system to see the actions only performed as a result of that playbook.
- **Status**: Status of the action. Available values are **All**, **Successful** and **Failed**.
- **Start/End Dates**: You can refine the action list by providing start and end dates of actions using the calendar buttons at both fields.
- **Action Value (Contains)**: A value to filter the action list where the action text contains this value.

There is a **Refresh** button on top right of the **Stage** field. You can click on this button to update the filtered actions list at that moment, or choose one of the predefined intervals in the button's dropdown list to update the list automatically at the selected interval.

There is also a **Download** button on top right of the list view. You can download your filtered action list as a CSV file to your computer using this button.

Enrichment History

Select **RESPOND** > **Status** > **Enrichment History**

Enrichment History tab lets you display and search logs of executed enrichments. To manage enrichment history, click on the **Enrichment History** tab in **Status** menu.

The page allows you to filter the action list by the following criteria as well as date:

- **All Integrations**: You can filter based on different integrations.
- **Submitters**: You can filter by users or automation.
- **Status**: Status of the enrichment. Available values are **All**, **Completed**, **Failed**, **Long Running**, **Not Started**, **In Progress** and **Excluded.**

There is a **Refresh** button on top right. For each entry there's also a **Result** column that will include a **Show** button to display the raw result of the enrichment.

Process Queues

Select **RESPOND** > **Status** > **Process Queues**.

Process Queues tab contains the following queue sub-tabs:

- **Alert Queue**: Lists the alerts received from any alert source that are saved in the SOAR database (including base events for applicable alert sources) and waiting to be processed (create/update Cases, execute playbooks).

  You can use the **Clear** button at the very end of queue list to clear the items in the respective queue.

- **ArcSight Listener Queue** : Lists the messages received from Micro Focus ArcSight ESM that are waiting to be processed and auto-enriched before they are added to the alert queue.

Managing Out Of The Box Workflows

Select **RESPOND > Playbooks > Workflow Templates**

> You must be an Administrator or a Superuser to create or import playbooks.

Out of the Box Playbooks provide the templates to help you design and implement your playbook. These templates are pre-designed workflows and provide guidance to customize automated response as per your requirements.

List of Out Of the Box Playbooks

ArcSight SOAR provides the following out of the box playbook templates:

- Access Attempts on Unidentified Protocols and Ports
- Admin Account Check
- Block Malicious IPs - CheckpointFW
- Block Malicious IPs - Palo Alto Panorama
- Check IP Reputation from Multiple Sources
- Command and Control Traffic-1
- Command and Control Traffic-2
- Command and Control Traffic-3
- Command and Control Traffic-4
- Endpoint Investigation - Windows
- Internal Scanning Device
- Multiple Authentication Failure
- Outbound Traffic to Suspicious Countries, Ports, Services
- Phishing Email
- Stolen-Lost Device
- Virus Traffic in the Network
- Investigate Suspicious User Account on OKTA
- APIVoid URL Enrichment
- Email Address Enrichment and Block on Cisco Ironport
- Email Address Enrichment and Block on FortiMail

- Email Address Enrichment and Block on Sophos XG
- Email Address Enrichment and Block on Symantec GW
- Investigate File Hashes & Block on Carbon Black
- Investigate File Hashes & Block on Checkpoint R80
- Investigate File Hashes & Block on Kaspersky SC
- Investigate File Hashes & Block on McAfee NSP
- Investigate File Hashes & Block on SEP Manager
- IP Enrichment on Free TI Databases
- URL Enrichment and Block on Check Point R80
- URL Enrichment and Block on McAfee Web GW
- URL Enrichment and Block on Palo Alto Panorama
- URL Enrichment and Block on Sophos XG

Prerequisites for Out of the Box Playbook:

To configure and use out of the box playbooks, a set of integrations/analyst tasks/lists, as listed in respective playbook guides, must be configured on your environment. You can also view the overview and prerequisites of each Out of the Box Playbook in the **Workflow Template** tab in the SOAR application.

Customizing Out of the Box Playbooks

The out of the box playbooks must be customized to create a playbook as per your requirement.

To customize out of the box playbooks:

1. Click **Workflow Template** tab.
2. Click **Create Workflow** and specify a name to the workflow in **Create Workflow From Template**window.
3. After importing the playbook as a template, select it and click **Repair** to configure as per your requirements.
4. Set parameter values as specified in the respective Playbook guide, in the **Workflow Repair Wizard** window.

Storing Lists in SOAR

Select **RESPOND** > **Configuration** > **Lists**.

The system can store a diverse set of values in lists. The lists are used as lookup tables for referencing purpose.

Searching a List

You can search a specific list, through the **Search** option. Click the button next to search, to view search results based on list's **Name, Content Type**and **Size** filters.

Creating a List

Click the **+Create List** button to create a new list. In the **List Editor** window, specify the details of list as follows:

1. **List Name**: Specify the name of the list that you want to create. For Example, VIP user names.

2. Select the option **Add Column** to create a new list. You can select Delete Column option if you are modifying an existing list.

3. Select the type of the data for the list from the drop down menu, for example, user name.

4. Specify the name of the column and click **Add**. The column name is displayed in below pane.

5. Enter a list item in the text field below the column name and then click **Add List Item** to add the list item for the newly created column.

   > Enter the list item in JSON format.

6. (optional) Enter a list item and click **Search** for searching a list item.

7. Click button to display the console pane. Click **Expand JSON**, to view list item in a JSON formatted order on the console.

8. Under **Actions** tab, click **update** to add the list item in the list or click **discard** to remove the list item.

9. (optional) Select the **Restrict Actions** checkbox to ensure that any action can not be taken (even if the action is a part of the playbook instructions) on the list items defined in the

newly created list.

10. (optional) Select the **Restrict Enrichments** checkbox to ensure that any enrichments can not be fetched (even if the fetching enrichment is a part of the playbook instructions) for the list items defined in the newly created list.

Example use case:

You might create a list to store the IP addresses of your data center.When you mark the list for **Restrict Actions** checkbox, SOAR will not take any actions for the servers listed in the list even if they are involved in a case. For example, your play book might contain a step to block all IP addresses on the Case scope, however it will not block those addresses defined in this list.

As another use case, you might define a list of VIP usernames. When you mark it for **Restrict Enrichments** checkbox, SOAR will not perform enrichments on these VIP users.

Editing, Deleting and Downloading a List

You can edit an existing list by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, **List Editor** window is displayed. Specify the details as per your requirement and click **Save**.

You can delete an existing list by clicking the **Delete** button under the **Actions** column.

You can also download a list as a text file (txt) by clicking the **Download** button.

Configuring Rest Clients

Select **RESPOND** > **Configuration** > **Rest Clients**.

To integrate with SOAR, a third party application needs a set of credentials generated at SOAR application. You can create these credentials at Rest Client.

When you click on the **Rest Clients** page, a list of case label is displayed.

Searching a Rest Client

You can search a specific third party application integration details, through the **Search** option. Click the ⚙▾ button next to search, to view search results based on **ID**, **Client ID**, Description, **Last Modified By,Modification Date** and **Actions**.

Creating a Rest Client

Click the **+Create Rest Client** button to create a new Rest Client. In the **Rest Client Editor** window, specify the required details in the **following** field and click **Save**.

Value	Description
Client ID	Specify the rest client ID.
Description	Specify the description of the rest client.

When you create a rest client by clicking save, a client secret is created for this rest client and displayed in the **Rest Client Details** window.

Ensure to note down the rest client secret along with the credentials as these would be needed whenever you call SOAR application using the REST API.

> **Note:**If you have lost the **Client ID** and **Client Secret** that you created for the rest client then you can not call the SOAR application using the respective REST API. In such cases, you must create the **REST Client** credentials along with the **Client Secret** again.

Editing and Deleting a Rest Client

You can modify an existing rest client by clicking the **Edit** button under the **Actions** column. When you click the **Edit** button, the **Rest Client Editor** window is displayed. Specify the description in the editor window as per your requirement and click **Save** to modify.

You can also remove an existing rest client by clicking the **Delete** button under the **Actions** column.

Managing Your Service Provider Contracts

To use this feature, you must have an MSSP contract with Micro Focus. Not available to customers in a SaaS environment

Select **ADMIN** > **Contract & Usage**.

Micro Focus provides a **pay-per-use program** for Managed Security Service Providers (MSSPs). This program offers our Partners a more affordable "pay as you go" option instead of maintaining of a perpetual license that requires a large initial investment.

Fusion helps you submit reports about daily and monthly average EPS (events per second) usage. You simply enable the MSSP feature, create an MSSP profile, and add contracts.

To get started, select **Add Contract**. To enable the MSSP feature, see the *Quick Start to Reporting EPS Usage*.

Managing Your MSSP Contracts

Not available to customers in a SaaS environment

Select **ADMIN** > **Contract & Usage** > **Contracts**.

Micro Focus provides a pay-per-use program for Managed Security Service Providers (MSSPs). This program offers our Partners a more affordable "pay as you go" option instead of maintaining a perpetual license that requires a large initial investment.

Understand MSSP Contracts

Select **ADMIN** > **Contracts & Usage** > **Contracts**.

A Partner subscribes to a Micro Focus MSSP contract to pay-per-use, with an entitlement of one or more of the ArcSight capabilities but without the initial cost of deploying the ArcSight Platform and Fusion in your IT environment. Each contract has a set length of time before it expires. You can access and use Fusion as long as you have a valid contract.

The MSSP contracts charge you basis EPS (events per second) usage. Micro Focus bases the fee on a tiered rate. The more daily average EPS you have, the less each event costs you. Monthly average EPS and cost do not impact the total cost but are an aggregate of the daily EPS and cost. If you have few events, Micro Focus charges you more for the service. You can see the tiers and the rates for each tier in ArcSight Platform and Fusion. The Tier Rates section lists the different tiers and the cost per EPS.

Add or Update a Contract

Select **ADMIN** > **Contract & Usage**.

After you purchase an MSSP contract from Micro Focus, you receive a copy of that contract. You must add the contract received to ArcSight Platform and Fusion you purchased. The contract also includes a signature file that Micro Focus uses to verify whether the contract you added is valid or not.

When you add the contract to ArcSight Platform and Fusion, Micro Focus verifies it and does not rely on the application's verification.

1. Click **Add Contract**.
2. Drag the contract and drop it in the box.

   or

   Click **Browse**, then browse to and select the contract.
3. (Optional) Click **Update** to modify an active contract.
4. (Optional) Click **Remove** to delete a pending future contract only.
5. Follow the instructions to complete the process.

   > If your contract expires or is going to, request Micro Focus for a new contract file.

Reviewing and Reporting EPS Usage

Not available to customers in a SaaS environment

Select **ADMIN** > **Contract & Usage** > **Monthly Usage Details**.

Fusion allows you to view your events per second (EPS) usage and the rates that Micro Focus charges in a single location. Fusion provides daily averages of your usage and an aggregate of your monthover-month usage and rates.

> ⚠️ It is possible that the bill you receive from Micro Focus might not match the cost displayed in the monthly report that you send to Micro Focus. The monthly report provides usage information, but Micro Focus could adjust the total cost, which results in a different invoiced amount. If you have any questions, please contact your Micro Focus representative.

Review Monthly Usage

Select **ADMIN** > **Contract & Usage** > **Monthly Usage Details**.

Fusion provides daily averages of your usage, an aggregate of your month-over-month usage and rates along with monthly usage reports and a yearly rate. The reports show the daily usages where the monthly value is an average of the daily use.

The monthly usage view provides usage or cost details. You can switch between the two views by clicking Usage or Cost in the upper-right corner.

> For various views available, the system rounds off usage and cost figures to two decimal places for easier on screen and print representation, but uses actuals in all calculations for accuracy.

1. (Optional) To view data for a different year, change the value for **Showing data for the year** at the top of the chart.

   By default, Fusion shows data for the current year.

2. Click **Usage** or **Cost** to view the daily average EPS usage or cost.

3. Review the usage or cost:

   Average Daily EPS (Aggregate)
   The Average Daily EPS (Aggregate) chart displays an average of your daily usage or cost. It then displays an aggregate (sum) of your month-over-month usage or the cost of your usage.

   Monthly Overview Table
   The Monthly Overview table displays your customers, their usage, their daily usage in a graph, their cost, and their month-over-month usage. You can download or email a PDF or CSV version of the report.

Submit an EPS Usage Report

Select **ADMIN** > **Contract & Usage** > **Monthly Usage Details**.

Fusion provides EPS usage reports that you can download or email in PDF or CSV format for your use. Each report is associated with a particular month of the selected year. It also allows you to automatically perform additional actions with these reports. You can:

- Create an email distribution list for the reports.
- Automatically send the monthly usage report to Micro Focus and the email distribution list.
- Set up a monthly reminder for sending the reports.

> To email PDF or CSV files from Fusion in a non-SaaS environment, your Arcsight Platform administrator must configure an SMTP server. If the server is not configured or available, you can still add the email addresses but Fusion will not be able to send the emails.

These reports contain a signature file that accompanies the email distribution or download of the PDF or CSV file to ensure that it has not been modified. Your browser might therefore prompt you to download more than one file when you attempt to as it includes the signature. When this occurs, agree to proceed.

1. In table row of the specific month that you want to report, select ….

2. Select the format for downloading or emailing the report:

   - **Download as CSV**

   - **Download as PDF**

   - **Email PDF**

   > Individuals on the email distribution list can ignore the signature file that accompanies the sent report.

Managing Your MSSP Profile

Not available to customers in a SaaS environment

Select **ADMIN** > **Contract & Usage** > **Profile**.

The MSSP Profile shows your account number and name from the contract, along with your contact information, while Fusion distributes the usage reports. You cannot access any of the sections of the MSSP Profile unless you add your contract to ArcSight Platform and Fusion. Also, to support emailing reports, your ArcSight Platform admin must set up an SMTP server.

Edit the MSSP Profile

Select **ADMIN** > **Contract & Usage** > **Profile**.

To be able to distribute the Fusion usage reports, you must set up your MSSP profile. Before you can set up your profile, you must add a contract. Your MSSP profile contains your account number and name from the contract, your contact information, and how you want to distribute the reports. The first time you log in, you must edit the partiality configured MSSP profile to have Fusion email the usage reports.

1. Click **Edit**.

2. Use the following information to set up or edit your profile:

   MSSP Identification
   Verify that the console displays your correct account information from the contract. You cannot edit these fields. If there is an issue, contact your sales representative.

   Contact
   Specify the name, title, phone number, and email address of the contact person if anyone has questions about the usage reports.

   Security Ops Center (SOC)
   Specify the SOC Identifier (name) and country of the SOC you are using for ArcSight Platform and Fusion.

3. (Optional) Configure how you want to distribute the usage reports.

4. Click **Save** to save your configuration information.

Configure Distribution of the Usage Reports

Select **ADMIN** > **Contract & Usage** > **Profile**.

The following information helps you configure how to distribute the Fusion usage reports. The Platform administrator must set up an SMTP server for the emails to work. You can add the email addresses without the SMTP server configured but Fusion does not send the emails, then.

- Send Reports Automatically
- Set Up a Monthly Reminder
- Create an Email Distribution List

For more information about setting up an SMTP server in a non-SaaS environment, see Connecting to Your SMTP Server in the *Administrator's Guide to ArcSight Platform*.

Send Reports Automatically

Fusion can automatically send the usage reports to Micro Focus and your email list. Fusion sends the reports on the first day of each month or on the day of first login, the same month.

1. Click **Edit**.
2. Under **Email Usage Reports** > **Email Settings**, enable **Automatically email usage reports**.
3. **Save** your changes.

Set Up a Monthly Reminder

If you chose not to automatically email usage reports, you can configure a reminder email to be sent to you so that you can send the email to Micro Focus and any other users.

1. Click **Edit**.
2. Under **Email Usage Reports** > **Email Settings**, enable **Remind me every month**.
3. **Save** your changes.

Create an Email Distribution List

Fusion allows you to create an email distribution list to make it easy to send the usage reports to the appropriate people.

1. Click **Edit**.
2. Under **Email Usage Reports** > **Email List**, click **Add Email Address**.
3. Specify the appropriate emails addresses.
4. **Save** your changes.

Managing and Importing Stored Data

*You must have the **Manage Storage Groups** permission to use this feature.*

Search performance can be affected by your environment's set up and the way that your data is organized. To enable faster search times, you can configure ArcSight to organize data into storage groups, which represent partitions in the ArcSight Database.

These storage groups can support compliance requirements for data retention policies, such as those for the Payment Card Industry Data Security Standard (PCI DSS). For example, you might be required to retain certain data for 12 to 24 months. You can instruct the ArcSight Database to purge data that is older than a certain number of months. By deleting data, you reduce the amount of content within the database and improve search performance.

Chapter 2: **Managing Your Stored Data**

You must have the Manage Storage Groups permission to use this feature.

Select **Configuration** > **Storage**.

The **Storage Information** list provides an overview of all available storage groups. You can have up to 10 storage groups, each with specific retention periods and query filters. To find a storage group, use the **Search** field.

Use Storage Groups to Organize and Retain Data

You can divide data into **storage groups**, which allows you to partition the incoming events data and provide different retention periods, based on the query filter. Because you can set data retention policies per storage group, you can retain certain high volume events for a short time period and other important events for longer time period. Higher volumes of event data, require more storage space. The **storage utilization** column displays the amount of storage utilized.

The **query filter** enables you to associate a storage group with specific compliance requirements, business needs, or search activities. Your specified query filters direct events to the correct storage group. For example, one group might have a filter for `categoryDeviceGroup =/ Firewall` and another for `severity >= 7`. If an event does not match any of the active filters, the event gets sent to the *Default Storage Group*. You cannot change the name, query, or rank of this built-in group.

> By default, the maximum value for retaining events in the *Default Storage Group* is 12 months. However, the license for your deployed product might require a lower maximum value, such as 30 days. For more information about how deployed products affect data retention policies in a non- SaaS environment, see " Understanding License Keys " in the *ArcSight Platform Administrator's Guide*.

The **Apply Changes to System** option at the top of the Storage Groups page indicates that one or more groups have been modified but the changes need to be applied.

- "Create a Storage Group" below
- "Direct Events to the Correct Storage Group" on the next page

Create a Storage Group

You can have up to **10 storage groups**, including the provided *Default Storage Group*.

1. Select **Configuration** > **Storage**.
2. Click the **add** icon **+**.
3. Enter a name for the storage group.

> ⚠ **CAUTION**: You cannot change the name after you create the group.

> The name cannot include special characters other than a hyphen (-).

4. Enter a query with which to filter the incoming events into this storage group.

   For example: `categoryDeviceGroup='/Firewall'` or `categoryDeviceGroup='/IDS'`.

   The query can include parentheses, quotes, and single quotes.

5. For the storage group's status, indicate whether to activate the group.

6. (Optional) For **Delete Data Older than**, enter the age of data, in months, that you want to purge from the storage group in the database.

7. Click **Save**.

8. Apply your changes.

Direct Events to the Correct Storage Group

For efficient data retrieval, the system matches each incoming event with the query filter for a single, active storage group. However, an event could be associated with the rules of more than one group. When an event matches with multiple storage groups, the system **assigns the event to the highest ranked group**.

For example, *if Event_29* matches the query filter for the storage groups ranked 3, 5, and 6, then the system assigns the event to the group that is ranked 3. If an event does not match any of the active filters, the system sends the event to the *Default Storage Group*.

You can change the ranking of storage groups to ensure that the system places events in the best location.

1. Select **Configuration** > **Storage**.

2. From the **Storage Information** table, drag each storage group up or down to the preferred priority position.

   The system always places the *Default Storage Group* in the lowest ranked position.

Activate and Deactivate Storage Groups

The application allows you to have up to **10 storage groups**, including the provided *Default Storage Group*. To deactivate to prevent new events from being sent to the group, change a storage group's status. For example, you might no longer need a particular storage group or find that you have changed the filters and functionality of that group from its original purpose. Rather than continuing to modify an existing group, you can deactivate it. Alternatively, you might want to activate a storage group only during certain periods of time.

Although you deactivate a group, the deletion settings for that group remain in effect.

1. Select **Configuration** > **Storage**.
2. Select the storage group that you want to activate or deactivate.
3. To edit the group's settings, click the   icon.
4. For **Group Status**, slide the indicator left or right.

   Activated groups display a status of **Active**.
5. Click **Save**.

Change the Settings of a Storage Group

After creating or modifying storage groups, you must apply the changes. You can modify multiple groups before applying your changes.

- "Modify a Storage Group" below
- "Apply Your Changes to a Storage Group" below

Modify a Storage Group

You can modify a storage group at any time.

1. Select **Configuration** > **Storage**.
2. Select the storage group that you want to modify.
3. Click the **pencil** icon 🖊.
4. For **Group Status**, slide the indicator left or right.
5. Activated groups will display a status of **Active**.
6. Click **Save**.
7. Apply your changes.

Apply Your Changes to a Storage Group

Select **Configuration** > **Storage** > **Apply Changes to System**.

When you change the query filter, status, or rank of a storage group, your changes do not go into effect until you apply the changes. The following considerations affect how your changes are applied:

- If you modify the query filter, the system will begin adding events that match the updated filter. However, the storage group retains all currently stored events associated with the previous filter. The retention policies continue to apply to all events within the group.
- If you do not want the storage group to have both sets of events, you can create a new storage group for the updated query filter, then deactivate the older storage group.

- On the first day of the month, the ArcSight Database deletes events matching the retention policies of the storage groups. For example, on March 15, you change the deletion time to three months from four months. On April 1, the database begins deleting all data older than three months.

- While changes are being applied, you cannot create or modify a storage group.

Use Storage Group Queries in a Search

Search allows you to include a storage group in a query. Rather than entering the query filter of a storage group again in Search, specify the following for your Search query: `Storage Group = Firewall Events`. By specifying the storage group, you limit the search to that storage group's partitions only, thus improving search performance.

Configure Retention Policies for Your Data

Events are stored in their assigned storage groups in the ArcSight database. Over time, the storage system can retain unneeded or outdated data. To preserve space in the database and improve data retrieval from storage groups, you can configure the database to remove events older than a certain number of months. For example, your data retention policy might expect your system to purge certain data, such as DNS logs that are older than 24 months.

When setting the policies for storage group retention and disk space utilization, do not allow your disk space utilization to increase above 90%. Running out of disk space can reduce the performance of searches due to increasing fragmentation. If such a situation continues to where there is no space left, then the database cannot ingest new data.

Delete Old Data from Storage Groups

Events are stored in their assigned storage groups in the ArcSight Database. Over time, the storage system can retain unneeded or outdated data. To preserve space in the database and improve data retrieval from storage groups, you can configure the database to remove events older than a certain number of months. For example, the data retention policy for your organization might expect data older than 24 months to be purged. This process **deletes data from the database**.

The system automatically applies all deletion settings on the first day of the month at 2:10 a.m.

1. Create or modify a storage group.
2. For **Delete Data Older Than**, enter the age of data, in months, when you want old events to be deleted.

   > By default, the maximum value for retaining events in the Default Storage Group is 12 months. However, the license for your deployed product might require a lower maximum value, such as one month. With a Log Management and Compliance license, you can choose **Never Expire** for a long-term storage option. To select that option, yrole must have the permission *Never Expire Search Results*.

   Ensure that your retention policy takes into consideration the maximum size of your storage groups and database. Also, consider that, in deleting events, the policy might affect results of an Event Integrity Check.

3. Click **Save**.
4. Apply your changes.

Manage Retention Policies for Imported Logger Data

Not available in a SaaS environment

> ⚠️ You can have up to 10 storage groups overall (including both groups created in the ArcSight Database and those migrated from Logger). Exceeding this quantity will likely affect the performance.

To manage the storage and expiration of the recently imported data from Logger, the system automatically enables the retention policy for the Logger event data. You do not need to manually direct events to a certain storage group and implement additional retention policies, but rather take advantage of the same storage rules already set in Logger.

You can update and review the retention policy strictly from the Logger interface. However, changes made to the retention policy after the archive migration has been started won't be reflected on the archives imported to the ArcSight Database.

> 🏠 For more information about retention policies for event data, see "Storage" in the *Administrator's Guide to ArcSight Logger*.

For a non-SaaS environment

As part of the metadata migration, the tool brings in the Storage Groups and their retention information. Because in Logger the retention settings are based on days, the tool converts them to months and rounds it up as to comply with the format used by the ArcSight Database. If the Maximum Archives Age value was set to -1 for a Storage Group in Logger (meaning a disabled retention policy), the retention process for it will also be disabled in Log Management and Compliance.

Importing Event Data From Logger

If you have ArcSight Logger deployed, you can import event data from Logger into the ArcSight Database to enable you to search on that data from within Fusion. To do so, the process requires you to first import the Logger event metadata, then Logger event archives into Fusion. After the system successfully imports a Logger archive file, the events in that archive file will be available in the Database. You can gradually import the Logger archive files as needed.

Chapter 2: **Prerequisites and Considerations for Importing Logger Data**

Does not apply in a SaaS environment

Since this process involves different ArcSight products interacting with each other, ensure that you have the correct credentials and requirements for all of them before you proceed.

- "Considerations for Importing Logger Data" below
- "Prerequisites for Logger" below
- "Prerequisites for the Fusion Capability" on the next page
- "Prerequisites for the ArcSight Database" on the next page

Considerations for Importing Logger Data

Please review the following considerations that affect how you can migrate data from Logger to the ArcSight Database.

- The process **imports only the archived events from the current Logger instance**. The process does not migrate content, configuration, and data from Logger peers.
- Logger event ingestion can continue up until the "Import Metadata for Logger Events" on page 340 step. It's advised to stop all Logger event ingestion, switch connectors to send events to ArcSight Database and then archive all the existing events in Logger before importing logger metadata
- You can perform only one migration at a time. If you plan to migrate data from different Loggers, run the migrations sequentially.

Prerequisites for Logger

- Admin user with SSH credentials.
- The username and password that you use to import Logger data must match the OS credentials set in Logger.

- The system directory must have enough space. For more information, see the *Release Notes for ArcSight Logger*.
- Whether the Logger host is installed in the same machine or a different one, it will need to have the **VSQL Client Driver** installed (this will be explained in Step 1 of the procedure).

Prerequisites for the Fusion Capability

- This procedure requires ArcSight Platform 22.1.2 at a minimum. See the Checklist for information about upgrading to the latest version.
- Admin user with ArcSight Database credentials.
- The system directory must have enough space. For more information, see the *Technical Requirements for the Arcsight Platform*.
- If the Fusion capability is deployed on the same machine as Logger, make sure that the RHEL/CentOS version used in Logger is supported by Fusion. For more information, see the *Release Notes for ArcSight Logger* and the *Technical Requirements for the Arcsight platform*.
- The Fusion capability must be reachable from the Logger instance (on port 5433).

Prerequisites for the ArcSight Database

- For the migration process, the user must have the *Logger Data Migration* permission assigned in Fusion. This is assigned by default to the *System Admin* role, but the user could have a custom role that includes the permission.
- For search execution after the data has been imported, users must have either the *Default Role* or a role with appropriate Search permissions.

Checklist: Migrating Logger Data

Does not apply in a SaaS environment

Use the following checklist to migrate event data from Logger for use in Search. You must perform the tasks in the listed order.

	Task	See
☐	1. (Conditional) If you installed ArcSight Platform using the `arcsight-platform-installer-22.1.0.16.zip` file or a previous version, upgrade to ArcSight Platform 22.1.2 at a minimum	"Upgrading to 22.1.2" in the *Administrator's Guide to ArcSight Platform 22.1* or upgrade to a more recent version if available
☐	2. Ensure that you have read the considerations and can comply with the prerequisites for importing Logger dat	"Prerequisites and Considerations for Importing Logger Data" on page 329
☐	3. Install VSQL Client Driver in Logger	"Install VSQL Client Driver" on the next page
☐	4. Archive all live data in Logger	"Archive Live Logger Data" on page 334
☐	5. Get the `loggerToReconConstructMounts.sh` script from ArcSight Database	"Obtain the Construct Mounts Instructions Script" on page 337
☐	6. Run the script to get archives mounting instructions	"Execute the Construct Mounts Instructions Script" on page 338
☐	7. Run the mounting instructions and the `logger_migration_preconfig.sh` script in the ArcSight Database	"Execute the Instructions in ArcSight Database" on page 339
☐	8. Import the Logger event metadata	"Import Metadata for Logger Events" on page 340
☐	9. (Conditional) If credentials were changed after importing metadata, update the username and password information, and then update the Logger registration	"(Conditional) Update the Logger Registration" on page 342
☐	10. Import the Logger event data that you want to search	"Import Logger Events" on page 343

Install VSQL Client Driver

Does not apply in a SaaS environment.

The Logger host requires the VSQL client to perform the data migration procedure. If client is not present yet, follow these steps to install the VSQL Client driver.

1. [Download the TAR version of the driver](#).

   > ⚠ This feature requires, at a minimum, version 11.1 of the ArcSight Database

   > ✔ Tip: Micro Focus recommends to use the same version for database server and TAR driver. Refer to the Technical Requirements for Arcsight Platform for details on the supported version.

2. To extract the TAR from the directory, run the following command:

   ```
   tar xvfz vertica-client-[version] [OS].tar.gz -C /
   ```

3. From your home directory, add the PATH:

   ```
   cd ~
   ```

4. Open the file:

   ```
   vi .bashrc
   ```

5. On the PATH variable located at the **/opt/vertica/bin** file, add the vsql path:

   ```
   export PATH-$ANT_HOME/bin:$JAVA_HOME/bin:$PATH:$P4_
   HOME/bin:/opt/vertica/bin
   ```

   If the PATH variable is not found, create it:

   ```
   PATH=$PATH:/opt/vertica/bin
   ```

6. Save the changes:

   ```
   :wq
   ```

7. Press **Enter**.

8. Refresh the .bashrc file:

   ```
   source .bashrc
   ```

9. To verify VSQL has been installed, run the following command:

```
vsql --version
```

Archive Live Logger Data

Does not apply in a SaaS environment.

> The steps listed in this procedure must be performed on your ArcSight Logger

All live data in Logger must be archived before you attempt the migration process.

- "Configure the Archive Storage Setting" below
- "Add an Event Archive" on the next page

Configure the Archive Storage Setting

Required only if you have not previously configured this setting

If you are using the Logger Appliance, create the NFS or CIFS mount point. For more information, see the Storage and Remote File System sections in Chapter 6 of the *Administrator's Guide to ArcSight Logger*. If you are using Software Logger and intend to use an NFS or CIFS mount point, ensure that the external storage point is mounted on the machine on which Logger is installed. For more information, see your system's operating system documentation.

1. Go to **Configuration** > **Storage** > **Archive Storage Settings**.
2. Specify a mount location and an archive path for each storage group. You can specify a different path for each storage group, thus enabling Logger to archive events to a different location for each storage group.

   You can configure settings for all storage groups on the **Archive Storage Settings** page even if you do not intend to archive all of them. Logger enables you to only save the storage group paths that have a mount configured and ignore the empty fields.

   - On Logger Appliances: Select (from the list box) a path in the Archive Path field appended to the path specified in the mount location. This location can be an NFS mount, CIFS mount, which is configured using the Logger user interface.

     For example, if the mount location you selected refers to the path `/opt/ARCHIVES`, and the archive directory in that location is `archivedir`, then specify `archivedir` in the **Archive Path** field.

   - In Logger Software, enter a complete path where the archive file will be written in the **Archive Path** field. This path could be a local directory or a mount point already

established on the Logger host.

> 🏠 Tip: On Software Loggers, the **Mount Location** field does not exist.

3. Click **Save**.

   If all fields are blank or without any changes, Logger will display the message *No changes have been made*. Otherwise, Logger will acknowledge the configuration with the message *Archive Storage Settings saved successfully*.

Add an Event Archive

1. Select **Configuration** > **Storage**.

2. Select **Event Archives**.

3. Click **Add**.

4. For **Name**, enter a meaningful name for the new **Event Archive**.

5. Specify the **Start** and **End** dates in the m/dd/yy format, where m is month number, dd is the day of the month (with a leading zero if necessary), and yy is the two-digit year number.

   When the **Start** and **End** dates are different, one archive file per storage group, for each specified day is created. For example, that will be the case when you specify the following **Start** and **End** dates:

   Start Date: 8/12/19

   End Date: 8/13/19

   > 🏠 Note: If a day's events have already been archived, you will not be able to archive them again. If you try to archive the same day's events twice, Logger will display a message with the already archived day or dates. If you are archiving a range of dates and some of them have been archived, the archive process will complete, skipping any days already archived, and a message will display the

   And, if you configure both storage groups—**Internal Event Storage Group** and **Default Storage Group**, four archive files will be created as a result of this archive operation—two files per storage group for the specified two days.

   The **Event Archives table** (in the **Event Archives** page) lists the archives by an alias in this format:

   ```
   <archive_name> [<yyyy-m-dd>] [<storage_group_name>]
   ```

6. Select the names of the storage groups that need to be included in the archive.

7. Click **Save** to start archiving events, or **Cancel** to quit.

Note: You can cancel an in-progress archive operation at any time using the Cancel link that displays on top of the Event Archives page.

If corruption cases have been detected before, please see the instructions for how to sanitize an Event Archive in Chapter 5 of the Administrator's Guide for ArcSight Logger.

Obtain the Construct Mounts Instructions Script

Does not apply in a SaaS environment.

To obtain the instructions for mounting the archives in the ArcSight Database, complete the following steps:

1. Navigate to the scripts folder in the ArcSight Database server, by default `/opt/arcsight-db-tools/scripts/`.

   This is where the `loggerToReconConstructMounts.sh` script is located.

2. To move the script to the Logger Server from which you want to import Logger Archive events, execute the following command:

   ```
   scp /opt/arcsight-db-tools/scripts/loggerToReconConstructMounts.sh root@<LOGGER
   IP>/opt/
   ```

Execute the Construct Mounts Instructions Script

Does not apply in a SaaS environment.

To generate instructions for the mounting of the data, complete the following steps:

> The output instructions are for guidance purposes only. They can be used as-are in Logger Appliances, but for Software Loggers, which can save data locally or externally, you must make sure that the path contained in the instructions corresponds to the NFS mount you created when configuring archive storage. See "Archive Live Logger Data" on page 334.

1. Give the execute right to the script that you just copied on the Logger Server:

    ```
    chmod +x ./loggerToReconConstructMounts.sh
    ```

2. Execute the script:

    ```
    ./loggerToReconConstructMounts.sh $<INSTALL LOGGER PATH>
    ```

3. The instructions generated will consist of the **mkdir** command to create a directory, and the **mount** command to perform the actual mounting, for example:

    ```
    Getting the instructions for /opt/mnt/ARCH-141-203
    mkdir -p /opt/LOGGER_15214141203/opt/mnt/ARCH-141-203
    mount -t nfs 15.214.129.238:/opt/shared|nfs4 /opt/LOGGER_15214141203/opt/mnt/ARCH-141-203
    ```

    These instructions will be generated for each of the mounts to be migrated.

Copy these instructions to execute them in ArcSight Database.

If the process fails to find archives that can be migrated, no instructions will be generated, and you will be notified by a UI message.

Execute the Instructions in ArcSight Database

Does not apply in a SaaS environment.

You must configure the ArcSight Database to receive the Logger migrated data.

> 🏠 The following instructions need to be run as **root** user, or a user with sudo credentials

1. To mount the archives on the ArcSight Database nodes, from your Linux command line, execute the commands that you copied or came up with during the procedure in "Execute the Construct Mounts Instructions Script" on the previous page.

2. Run the `logger_migration_preconfig.sh` script located by default in the `/opt/arcsight-db-tools/scripts/` directory.

Chapter 2: Import Metadata for Logger Events

Does not apply in a SaaS environment.

Select **Configuration > Import Logger Data > Logger Metadata Import**.

> ⚠️ **This topic applies only to Logger processes soon-to-be shut down.**

Logger metadata refers to the information that is stored in the Logger `postgresql` database, which is needed to read the events from the Logger archive files for each storage group.

You import the metadata once for each Logger whose processes are soon to be shutdown. Complete the following activities:

- "Register a Logger " below
- "Import the Metadata" on the next page
- "(Conditional) Update the Logger Registration" on page 342

Register a Logger

Applies only if you have not previously registered the Logger from which you will import data

Before importing the metadata, make sure to add the Logger details for the import process.

1. In Fusion, select **Configuration > Import Logger Data > Logger Metadata Import**.
2. Click the **+** icon.
3. Add the Logger details such as:
   a. **Host**: Logger IP address or host name
      For example, `12.345.67.890` or `logger6.extremelyfocused.com`
   b. **Host Username**: OS username
   c. **Host Password**: OS password
4. Click **Save**. Otherwise, click **Cancel**.

> 🏠 **Note:** You can remove Logger registration if no data has been imported. To delete the Logger registration, click the delete icon (trash can).

Import the Metadata

> **Note:** It's recommended that you perform the following steps before the actual metadata import:
>
> - Stop all Logger event ingestion
> - Switch connectors to send events to the ArcSight Database
> - Archive all the existing events in Logger before importing the Logger metadata

While importing the metadata, the Logger server must be accessible at all times.

The metadata contains all the information related to accessing the events of a particular Logger. You can migrate the Logger metadata to the ArcSight Database directly from the **Logger Metadata Import** page.

> ⚠ Make sure to import the metadata before importing the Logger data as this is the first step to view and consume logger events.

1. In Fusion, select **Configuration > Import Logger Data > Logger Metadata Import**.
2. Check the box next to the Logger whose metadata will be migrated and click the **import** icon.

   A pop-up window will notify you that the Logger metadata import procedure is about to begin, making sure you have already mounted the appropriate archives on all database nodes.

   Here a determination must be made regarding what happens once the import of metadata is done. Choosing:

   - **Yes** on the pop-up window, the Logger processes will resume after the import is finished. The system proceeds to import and store the metadata in Log Management and Compliance.

   - **No** on the pop-up window, the Logger processes will remain shut down. The system proceeds to import and store the metadata in Log Management and Compliance.

     > **Note:** After the metadata has been imported, and the Logger processes shut down, Logger can be uninstalled and the appliance box can be re-purposed.

   - **Cancel** on the pop-up window, the metadata migration process will not be initiated.

(Conditional) Update the Logger Registration

If the credentials have been changed after registering a Logger, make sure to update the username and password information before importing the Logger metadata.

The Logger processes status, host username, and password can be updated after the Logger registration, but only if the metadata import process hasn't started.

These values cannot be updated after you start an import.

1. In Fusion, select **Configuration > Import Logger Data > Logger Data Import**.
2. Check the box next to the Logger host and click the pencil icon.
3. Update the values accordingly.

   Ensure that the username and password that you use match the OS credentials set in Logger.
4. Click **Save**. Otherwise, click **Cancel**.

Chapter 2: Import Logger Events

Does not apply in a SaaS environment.

Select **Configuration** > **Import Logger Data** > **Logger Data Import**.

This option will allow you to bring events from a Logger instance to the Fusion capability and perform searches on them. Since this process consumes both time and resources, consider importing only events in necessary time ranges.

> ⚠️ Before you can migrate Logger data, you must import the metadata that defines it.

- "Import Archived Events" below
- "Review Migration Details" on the next page
- "Resume an Incomplete Migration" on page 345
- "Delete Incomplete or Failed Migrations" on page 346

Import Archived Events

Before importing Archived Events, ensure that you comply with the prerequisites for the process.

1. Select **Configuration** > **Import Logger Data** > **Logger Data Import**.
2. Click **+**.
3. Select the Logger host of your preference.

   You can choose only one host at a time.
4. Specify the time range that you want to import, following these considerations:

   - The time range is based on receipt time.

     > Convert the time range you wish to search through from browser time/selected time zone to UTC.
     >
     > That way, once the data is imported, you can search through it using the original browser time/selected time zone.

   - The migration only allows you to migrate a minimum time range of 1 day.

- Specify a date in the past. You cannot import data for future dates as it will import no events and will cause issues when you try to import new data again.

- Overlapping dates will cause an error message. If this is not the first import of this Logger instance, ensure to select a time range different than the one already imported.

> Select a data-time range different than the one already imported. To confirm the host's start and end dates already available in the ArcSight Database, see how to verify the migration table in "Review Migration Details" below

5. Click **Import**.

6. To check the import progress, view the **Import Status** column.

   The import will take a considerable amount of time, based on the quantity of events that are present in the time range selected.

7. (Optional) If the import is interrupted, you can attempt to resume the process.

   Alternatively, you can delete an incomplete migration.

Review Migration Details

The migrations table will display the most relevant information of all the imports executed. For each migration, the system registers the following details:

Logger Host
    Represents the Logger IP address or host name. For example, `12.345.67.890` or `logger6.extremelyfocused.com`.

Data Start Date
    Indicates the absolute date of the earliest possible event.

Data End Date
    Indicates the absolute date of the latest possible event.

Import Date
    Indicates the migration date and time displayed in the ArcSight Database timezone.

Import Status
    Indicates the status of the import process:

- **Initialized**: The verification of the archives corresponding to the requested time range is being performed.

- **In progress**: Import is still in progress. Archived events are being extracted, read & sent to the ArcSight Database.

- **Complete**: Successful import execution.

- **Failed**: The archives are inaccessible, which can be caused by:

  - An unresponsive mount

  - A network connectivity issue

  - A user who doesn't have the correct access permissions

  - Data that couldn't be uncompressed, etc

  Ensure that you comply with the prerequisites before importing data.

Event Count
Indicates the number of events migrated. This number increases automatically as the process continues.

Logger Host User Name
Indicates the OS username associated with the Logger host.

Data Import ID
Represents the unique identifier for the event migration. You must have this value to delete a migration.

To review details about the executed migration, see the logs in the `opt/vertica/udfs/datamigration/logs/` directory.

After events have been imported, either Logger or the Fusion capablity will manage the retention policy depending on the state of the Logger processes.

Resume an Incomplete Migration

A migration might be interrupted if access to the mount or data file is affected in any way during the process: an unresponsive mount, a network connectivity issue, a user who doesn't have the correct access permissions, data that couldn't be uncompressed, etc.

An **Incomplete** migration can be resumed. The process starts from the last point of migration so you do not lose the data previously migrated.

1. Select the migrations that you want to resume.

2. Click ⟲.

A migration that continues to appear as **incomplete** after it has been resumed at least once, might indicate the data cannot be migrated because of corruption issues.

Check the logs for any related messages, and contact support to help finish the migration.

Delete Incomplete or Failed Migrations

It's possible that a migration might fail to complete. For example, the status is **Failed** or indicates that the migration is **Complete** but it contains no events. In these types of scenarios, you can delete the migration, then try again.

1. Select the migrations that you want to delete.
2. Click 🗑.

Creating and Using ArcSight Dashboards

Available only with ArcSight capabilities.

Select **Dashboard**.

The **Dashboard** enables you to visualize, identify, and analyze potential threats by incorporating intelligence from the multiple layers of security sources that might be deployed in your security environment:

- Managing and monitoring ArcSight infrastructure components with ArcSight Management Center (ArcMC)
- Real-time event monitoring and correlation with data from ArcSight Enterprise Security Manager (ESM)
- Analyzing end-user behavior with ArcSight Intelligence
- Performing deep-dive investigations with Log Management and Compliance
- Responding to and mitigating cyber attacks with ArcSight SOAR

To help you get started, Fusion provides a set of out-of-the-box widgets and dashboards. Users can organize the widgets into personalized dashboards. Out-of-the-box, any user can perform the following actions:

- View dashboards owned by or shared with the user
- Modify, delete, and export dashboards owned by the user
- Create or clone dashboards
- Import dashboards
- Set a dashboard as a personal default dashboard

You can create one or more ArcSight dashboards that incorporate widgets in your preferred arrangement. Depending on your role, you can create dashboards to be shared with specific roles, and even identify which of those dashboards should be the default landing page for a role.

Viewing a Dashboard

Select **Dashboard**.

The Dashboard automatically displays your default dashboard when you log in or select **Dashboard**. If you do not have a default dashboard, the Dashboard displays the list of available dashboards.

While viewing a dashboard, you can modify its settings or clone it to create a new dashboard.

View Data in a Dashboard

Select **Dashboard**.

Content in a dashboard depends on the widgets that it displays, as well as the dashboard's specified time range.

View a Different Dashboard

When viewing a dashboard, select **View All Dashboards**.

In the course of your day, you might need to switch among several dashboards. You can view the list of dashboards in two ways:

- "Favorite Dashboards" below
- "All Available Dashboards" below

The list indicates whether a dashboard is shared, for your personal use, or assigned as the default for a role. You can also see who owns each dashboard. An "out-of-the-box" label indicates that the dashboard is provided with the Dashboard. In general, out-of-the-box dashboards are available only to the Dashboard administrator because they require configuration before use.

Favorite Dashboards

You can specify which dashboards are your favorites.

All Available Dashboards

You can view the full list of available dashboards. A star beside the name indicates that you have marked that dashboard as a favorite.

Viewing Analyst and Entity Details

Some of the widgets in the dashboard allow you to review activity associated with specific cases, case owners or owner groups, and entities.

Case Overview by Owner

Select an owner in a widget.

You can review all cases currently assigned to a specific owner. When you select an owner in a widget, the Dashboard opens the **Case Overview by Owner** page. For each case, the table includes the following details:

- Severity of the case
- Current stage of the case
- Length of time that the case has been assigned to the owner
- Time since the case was created
- Time since the case was last updated

To determine when the owner received a particular case, hover over the **Owned** field. If you hover over the **Created** and **Last Updated** fields, the Dashboard shows the specific date and time that the case was created or last updated, respectively.

Review Entities

Requires ArcSight Intelligence be deployed. This feature is not available in a SaaS environment.

Select an entity in a widget.

You can explore the entities and their risky behaviors in the following ways:

- When you select an entity type in the Entity Count Overview widget, the Entities page opens in the ArcSight Intelligence UI, where you can view the details of the risky entities of the selected entity type. You can also navigate to the Entities page in the ArcSight Intelligence UI in the following ways:

    - When you have deployed only ArcSight Intelligence, click **ENTITIES AT RISK** in the left pane.

    - (Non-SaaS only) When you have deployed ArcSight Intelligence and Recon, click **INSIGHTS** > **Entities At Risk** in the left pane.

- When you select an entity in the Top Risky Entities widget, the Explore page opens in the ArcSight Intelligence UI, where you can explore the risky activities associated with the entity.

Managing Dashboards and Content

Select **Dashboard**.

You can add, remove, and rearrange the order of widgets in a dashboard. You can also change the content of a widget then save it with a unique name. To edit a dashboard, you must be currently viewing it.

Understand the Provided Dashboards

To help you get started, the Dashboard provides out-of-the-box dashboards with associated widgets. You will need to configure the widgets to ensure the dashboards display data appropriately for your environment.

- "Data Processing Monitoring" below
- "Health and Performance Monitoring" below
- "How is My SOC Running?" on the next page
- "Entity Priority" on the next page
- "Entity Risk " on the next page

Initially, the out-of-the-box dashboards are available to the administrative user created during the initial log in. This user can share these dashboards with SOC team members, who can then create their own clones for custom dashboards. Alternatively, administrators can create one or more clones based on these dashboards, then share the clones, and set default dashboards for roles.

Data Processing Monitoring

Requires that at least one deployed capability uses the ArcSight Database.

The dashboard, **Data Processing Monitoring**, provides information to monitor the rate of event ingestion into the ArcSight Database. It includes the following widget:

- Database Event Ingestion Timeline

Health and Performance Monitoring

Requires that at least one deployed capability uses the ArcSight Database.

The dashboard, **Health and Performance Monitoring**, provides information about the status of the ArcSight Database. It includes the following widgets:

- Database Cluster Node Status
- Database Event Ingestion Timeline
- Database Storage Utilization

How is My SOC Running?

Requires ArcSight ESM Command Center be deployed. This dashboard is not available in a SaaS environment.

The dashboard, **How is my SOC running?**, gives you an overview of the status and trends related to ESM case management. It includes the following widgets:

- Case Breakdown
- Case Load
- Case Timeline
- Case Workflow Analysis
- Productivity
- Threat Analysis Funnel

Entity Priority

Requires ArcSight Layered Analytics capability be deployed. This dashboard is not available in a SaaS environment.

The dashboard, **Entity Priority**, combines content from ArcSight ESM and Intelligence to provide the status of users and entities at risk. It includes the following widgets:

- Active List
- Entity Count Overview

Entity Risk

Requires ArcSight Intelligence be deployed. This dashboard is not available in a SaaS environment.

The dashboard, **Entity Risk** provides at-a-glance actionable information on the current, overall risk of your organization. It includes the following widgets:

- Analytics Pipeline
- Entity Count Overview

- Overall Risk Level
- Top Risky Entities

The dashboard provides the following information:

- Risk statistics: number of events analyzed, number of anomalies and violations found, and the number of active risky entities.

- The types of entities involved and their risk counts. When you click an entity type, the Entities page opens in the ArcSight Intelligence UI, where additional information for the selected entity type is displayed.

- The trending risk of the organization.

- The dominant potential threat, if any.

- The top 5 risky users. When you click a user, the Explore page opens in the ArcSight Intelligence UI, with the selected user's name applied to the anomalies and violations filter.

- An option to download a PDF containing a detailed report of the risk of the organization. For more information about PDF report, see the "PDF Reports" section in the *ArcSight Intelligence User's Guide*.

Change the Time Range of Data in a Dashboard

When viewing a dashboard, select 🖉.

Most of the widgets in a dashboard display data according to the either a specified **Time range** or an **As of now** setting, which displays data based on the last time that you refreshed the browser. You can configure the time setting.

If you select a preset time, the Dashboard displays data starting from 12:00:00 a.m. of the first date in the range to 11:59:59 p.m. of the last date in the range. If the last date is the current date, then the Dashboard defaults to the current time or time of the last browser refresh. For example, the **Last 1 month** setting might be from 12:00:00 a.m. April 29 to 3:34 p.m. May 29. Note that the Dashboard does not display minutes and hours.

To display time values, the Dashboard uses your browser settings, such as your local time zone.

Create or Clone a Dashboard

You can build as many dashboards that you need either by creating a new dashboard or copying a custom or out-of-the-box dashboard.

- "Create a Dashboard" below
- "Clone a Dashboard" below

Create a Dashboard

You can create as many dashboards as you need.

1. (Conditional) From within an existing dashboard, select ... > **Create new Dashboard**.
2. (Conditional) From the Dashboards list, select **+**.
3. Specify a **Title** for the new Dashboard.

   The title can be a maximum of 150 characters, and must be unique.
4. To add a widget, select **+** beside **Main Context**.
5. Choose the widget that you want to add.
6. Modify the widget's properties.
7. Continue to add widgets as needed.
8. Arrange the widgets how you prefer.
9. **Save** your changes.

   Alternatively, you might choose to clone an existing dashboard or import a dashboard that someone else created.

Clone a Dashboard

To quickly create dashboards, you can copy an existing dashboard. For example, Inez Bates wants to customize an out-of-the-box dashboard and share it with her APJ analyst team. She clones the dashboard, then modifies some of the widgets to include only cases that the team owns.

By default, the Dashboard copies the name of the original version and adds "Copy of" to the name. You can change that title as part of the cloning process or edit the title later.

1. From within an existing dashboard, select … > **Clone**.

2. Specify a unique name for the new dashboard.

3. (Optional) Indicate that you want to add the new dashboard to your Favorites.

4. **Save** your changes.

   Alternatively, you can import a dashboard that someone else created.

Modify a Dashboard

While viewing a dashboard, select 🖉.

You can only change the configuration of the dashboard that you are currently viewing, such as editing a widget's properties or adding and removing widgets.

- "Add Widgets" below
- "Modify a Widget's Properties" below
- "Rearrange the Order of Widgets" below
- "Remove Widgets" on the next page
- "Change the Dashboard's Name" on the next page

Add Widgets

While viewing a dashboard, select 🖉, then + in Main Context.

To find an existing widget, you can search by its name or the tags assigned to it. After choosing the widget, you can change its properties to suit your dashboard.

To group widgets in sections under the **Main Context**, select **Nested Context** from the widget selector or select a context that has already been added to the dashboard. Then you can add widgets in that section. You can also change the name of the sections.

Modify a Widget's Properties

While viewing a dashboard, select 🖉.

To edit the settings of a widget, select the widget. Make your changes in the **Widget Properties** pane. Then save your changes.

Rearrange the Order of Widgets

While viewing a dashboard, select 🖉.

To rearrange the order of widgets in a dashboard, simply drag each widget to the new location. Then save your changes.

Remove Widgets

While viewing a dashboard, select ✏️.

To remove a widget, select X within the widget's boundaries. Then save your changes to the dashboard.

Change the Dashboard's Name

While viewing a dashboard, select ✏️.

The title of a dashboard can be a maximum of 150 characters, and must be unique.

Mark a Dashboard as a Favorite

To more quickly find a dashboard, you can add it to your **Favorites** list.

While viewing a dashboard, select ☆.

Specify a Default Dashboard

Select … > **Set as default for me**.

When you log in, the Dashboard automatically displays the default dashboard that you have chosen or that an Administrator has assigned for your role. If no dashboard has been assigned to you or no default exists, you will see the list of available dashboards.

To override the default dashboard assigned to your role, you can specify any currently displayed dashboard as your preferred landing page.

Share a Dashboard

*You must have the **Share Dashboard** permission to perform this function.*

Select ... > **Share**.

You can share the currently displayed dashboard with one or more of your assigned roles. If you have the **Manage Roles** permission, you can share the dashboard with any role.

Alternatively, if you cannot share a dashboard, you can export the dashboard for others to import and use.

> You cannot re-share a dashboard that has been shared with you.

Import and Export a Dashboard

As an alternative to sharing or copying a dashboard, you can export the dashboard as a json file for other users to import to their Dashboard. The json file contains information about the dashboard's configuration and the included widgets. The file does not contain any data displayed in the dashboard. You can modify the exported json file or edit the imported dashboard.

For example, Inez Bates on the APJ analyst team really likes a dashboard that **Murphy Buckley,** on the EMEA team, made for his personal use. Murphy could share this dashboard with Inez. However, the widgets are configured for the AMS team's use, so the data would not be useful for Inez. Instead, Murphy exports the dashboard and sends the json file to Inez. She imports the dashboard, then modifies some of the widgets to point to cases that she and the APJ team own.

- "Considerations for Importing a Dashboard" below
- "Import a Dashboard" on the next page
- "Export a Dashboard" on the next page

Considerations for Importing a Dashboard

Changing the json file of a dashboard can cause problems either during import or within the Dashboard. Usually, you only need to change the name of the dashboard in the file. Before importing a dashboard, please review the following considerations:

- You cannot import a dashboard whose name already exists in your Dashboard environment. Ensure that you change the title of the dashboard in the json file.

  > This caveat includes names of dashboards that other users have created and which you might not see in your list.

- You cannot import a dashboard if it contains widgets that do not exist in your Dashboard environment.

Import a Dashboard

When viewing the list of Dashboards, select … > **Import Dashboard**. Then browse to the appropriate json file.

Export a Dashboard

When viewing a Dashboard, select … > **Export Dashboard**.

Display a Dashboard on the SOC Screen

Like most software, the Dashboard will end a session that has been idle for a while. This is good for security. However, it can be inconvenient if you display a dashboard on the large monitors in your SOC. To avoid manually interacting with the browser or logging in regularly, you can use a plug-in that automatically refreshes all content in the browser tab that displays the dashboard.

To automatically refresh dashboards on the SOC screen:

1. Install an Auto Refresh add-on for your browser.

   There are free add-ons available for supported browsers.

2. Specify the time interval after which you want the browser tab to refresh automatically.

   For instance, if you set the time for auto-refresh to five minutes, your browser tab will refresh automatically after an interval of five minutes.

3. (Optional) Minimize the left navigation pane.

Note that, when you refresh the tab, the Dashboard always updates to the latest data based on your chosen time range.

Configuring Widgets

Widgets display data according to your specifications. You can filter content by specific case owners or groups, case severities, and sub-filters.

Understand Widget Properties

When you configure a widget, you might see a combination of some or all of the following properties:

Title and Subtitle

Specifies the name and an optional secondary name for a widget you want to add to your dashboard.

You can also specify whether the dashboard displays the title or subtitle.

In general, because you might have several variations of some widgets, it's a good practice to title each widget according to your sub-filter criteria. For example, SOC Manager Franz Tupper creates a Case Breakdown widget for each of the SOC's three owner groups: EMEA, AMS, and APJ. He names the widgets *Case Breakdown-EMEA*, *Case Breakdown-AMS*, and *Case Breakdown-APJ*.

Severity

Specifies the categories of importance, or severity, assigned to the affected cases. For example, in ESM, some cases might be categorized as *Catastrophic* or *Marginal*.

When selected for **Group by** or **Facet**, you can add sub-filters by specifying the type of Cases, Assigned Owners, or Assigned Owner Groups that you also want to view.

Assigned Owners

Indicates that you want to display data based on the individuals assigned to the affected cases. You can specify the **Owners** that you want to include.

If you do not specify an owner, the Dashboard includes data for all owners. If you specify more than five owners, the Dashboard displays data for the top five selected owners. Then adds an **Other** category that totals the values for all other selected owners.

When selected for **Group by**, you can add sub-filters by specifying the type of Cases and Importance categories that you also want to view.

Assigned Owner Groups

Indicates that you want to display data based on the owner groups, or teams, assigned to the affected cases. The widget also displays all cases assigned to the individuals and child groups within the owner groups. You can specify the **Owner Groups** that you want to include.

If you do not specify an owner group, the Dashboard includes data for all groups, and thus all owners. If you specify more than five owner groups, the Dashboard displays data for the top five selected groups. Then adds an **Other** category that totals the values for all other

selected owner groups.

When selected for **Group by**, you can add sub-filters by specifying the type of Cases and Severity categories that you also want to view.

Assigned Cases

Applies only when you specify Severity for Group by

Indicates whether a sub-filter includes cases assigned to the specified owners.

To include specific owners or owner groups, select Owners then add the names that you want to include. Otherwise, the Dashboard displays data for all assigned cases.

In general, to view sub-filter data, you might hover over the visual in the widget or drill down into the data.

Unassigned Cases

Applies only when you specify Severity for Group by

Indicates whether a sub-filter includes unassigned cases.

Number of Groups

Applies only to the Productivity (SOAR) widget

Indicates whether a sub-filter includes the most productive number of groups.

Statuses

Applies only when you specify Statuses for Facet

Indicates whether a sub-filter includes statuses.

Show Top N Playbooks

Applies only to the Productivity (SOAR) widget

Indicates whether a sub-filter includes the number of Top Playbooks executed.

Classifications

Applies only to the Productivity (SOAR) widget

Indicates whether a sub-filter includes the classification of the attack type.

Number of Playbooks

Applies only to the Productivity (SOAR) widget

Indicates whether a sub-filter includes the number of Playbooks executed.

Target for Case Closure

Applies only to the Productivity and Case Load widgets.

Specifies the number of cases per week that you expect each owner group (Productivity widget) or owner (Case Load) to close.

Time Range

Specifies the start and end dates for the data that you want to view:

- **Dashboard's default** tells the widget to use the time range set for the dashboard.

- **As of now** tells the widget to use the most recent data retrieved from the data source.

  Data updates each time you refresh the browser, unless you have specified a **Custom** time range.

  > You can set a **maximum time range** to limit the amount of data that the Dashboard can collect from its data sources. For example, you can specify 365 days of data. For more information, see the *Administrator's Guide to ArcSight Command Center for ESM*.

To assign or change the severity or owner of a case, use the ArcSight Console or Command Center.

Layout

Specifies the orientation of the widget in a custom dashboard. For example, you might want the *Database Event Ingestion Timeline* widget to span the width of the dashboard.

Understand the Provided Widgets

The Dashboard ships with several widgets designed to help you manage your security operations. When you create or modify a dashboard, you can choose from the full set of widgets and configure them as needed.

The Dashboard provides the following out-of-the-box widgets:

Active List

Requires ArcSight Intelligence and ArcSight ESM be deployed for best effect. This widget is not available in a SaaS environment.

To watch for suspicious activity associated with entities, add **Active List** widgets to your dashboard. Each widget displays the top five at-risk entities, based on the specified **Active list**, **Field**, and **Entity type** settings with both ESM and ArcSight Intelligence installed.

The available active lists correspond to active lists in ESM. For example, you might have watch lists for privileged or administrative users or vulnerable hosts. If an active list entry matches an entity in ArcSight Intelligence, then the widget also shows the ArcSight Intelligence risk score for that entry. However, if the ArcSight Intelligence capability is not deployed, the widget cannot display risk scores but just entities in alphabetical order.

Analytics Pipeline

Requires ArcSight Intelligence be deployed. This widget is not available in a SaaS environment.

The **Analytics Pipeline** widget provides the risk statistics for the last analytics run. It displays the number of events analyzed, the number of anomalies and violations found, and the number of active risky entities. This widget also provides the option of downloading a PDF report detailing the current risk of the organization. You can select the orientation of the widget as **Landscape** or **Portrait**. The default orientation is **Landscape**.

Case Breakdown

Requires ArcSight ESM be deployed. This widget is not available in a SaaS environment.

The **Case Breakdown** widget displays the number or percentage of cases by their **Severity**, **Owners**, or **Owner Groups**. The widget always shows data **As of Now**, regardless of the specified time range for the dashboard.

By default, the widget shows data for total open, assigned cases. The widget displays a maximum of six data points, which comprise the top five objects associated with the specified filter plus an *Other* object that combines the rest of the cases. For example, if you have seven case owners, the widget shows specific values for the five owners with the largest quantity of cases, then groups the total number of cases for the other two owners in the Other category.

You can change the widget's properties to view cases in a different state, such as cases created by specific analysts. For example, SOC Manager Franz Tupper wants to view all cases created by his Level 1 analysts. He sets the filter to **Assigned Owners**, and in the sub-filters specifies Jin Stafford, Neve Marshall, Troy Leach, and Chole Gay as **Owners**. Then he selects **Created** for the state that he wants to analyze. The widget will display the quantity and percentage of cases created by each analyst. Because Franz has configured the dashboard to automatically refresh, he sees in real-time when the analysts add new cases.

If you don't specify an owner or owner group, the widget displays data for all cases.

Case Load

Requires ArcSight ESM be deployed. This widget is not available in a SaaS environment.

To help managers balance the amount of work assigned to case owners, the **Case Load** widget provides several case management metrics:

- Average number of cases each owner closes per week
- Estimation of the time required to close all cases currently assigned to the owner based on the time elapsed since the cases were opened
- Projection of the number of cases per severity that the owner might not be able to close, based on the configured target, the time elapsed since the cases were opened, and the average velocity of the owner. This assumes that owners work on cases in severity order, from highest to lowest.

By default, the widget shows the data for total open, assigned cases for the top three members of the group based on their average number of cases per week. You can filter the data by specific **Owner Groups**. The metrics are based on the specified time range and the target number of cases that you expect the owners to close per **Severity**

For best use of this widget, it is recommended that you create one Case Load widget per owner group. In this way, you will see details for members of the owner group.

Case Timeline

Requires ArcSight ESM be deployed. This widget is not available in a SaaS environment.

The **Case Timeline** widget shows changes in the volume of cases over a specified time range. By default, the widget filters the data according to the **Severity** category assigned in ESM. However, you can also choose to view trends for other case states, such as cases **Closed** by specific **Owners** or **Owner Groups**.

To observe the breakdown of cases associated with a specific date, you can hover over any location within the timeline. You can also zoom in to view a particular time range, either using the magnifier icons or by clicking and dragging within the graph.

Case Workflow Analysis

Requires ArcSight ESM be deployed. This widget is not available in a SaaS environment.

The **Case Workflow Analysis** widget helps you compare the current volume of cases per stage with how the cases transitioned among the stages. In the widget, the width of the lines indicates the average time cases have taken to move from stage to stage during the specified time range. The diameter of each circle, except for the *Closed* stage, represents the total number of cases currently at that stage, based on the last refresh of data from the source.

> The widget does not represent backward transitions. For example, a case moves from *Final* back to *Follow-up* during the specified time range.

By default, the widget shows data for total open, assigned cases. You can also choose to filter the data by **Severity**, **Owners**, or **Owner Groups**.

Database Cluster Node Status

Requires that at least one deployed capability uses the ArcSight Database.

The **Database Cluster Node Status** widget helps SOC managers and IT administrators monitor the state of the nodes that host the database. This widget displays the state of each node in the database cluster. It also raises awareness that the number of nodes that are down can affect the resiliency of the database cluster. For example, if the database resiliency setting is 1, and two of three nodes go down, then the database might automatically shut down to protect itself.

Also, when nodes are down or recovering from a failure, it's possible that you might experience data loss. The longer that a node is offline, the longer it will take to recover because it needs to acquire the data available in the rest of the cluster.

Database Event Ingestion Timeline

Requires that at least one deployed capability uses the ArcSight Database.

To help SOC managers and IT administrators monitor the rate of event ingestion into the database, use the **Database Event Ingestion Timeline** widget. Due to differences in how quickly an event from different sources arrives at the database for storage, the moment when a database stores an event differs from when the event occurred. This widget measures when the database receives the event data.

Database Storage Utilization

Requires that at least one deployed capability uses the ArcSight Database.

To help SOC Managers and IT Administrators ensure that disk use does not overload the database nodes, the **Database Storage Utilization** widget displays storage utilization data for up to five database nodes. In general, most administrators keep disk usage below 60 percent per node, thus ensuring space for temporary activity required by some query execution operators.

If the database cluster has more than five nodes in the cluster, you might specify the nodes with the least amount of free space available. In this way, you can monitor the nodes at most risk of running out space. For each node, you can compare the percent and quantity of space used to the total amount. You can also monitor the throughput and latency of the database per second.

The ArcSight Database supports use of a third party storage location technology, shared among its database nodes on premises or cloud. This shared storage location is also called Communal Storage and represented in the associated widget.

> The computational and communal layers of the database are separate and allows storage of data in a single location with the ability to elastically vary the connected computer nodes per necessary computational needs. For more information, see the *Administrator's Guide to ArcSight Platform*.

Entity Count Overview

Requires ArcSight Intelligence be deployed. This widget is not available in a SaaS environment.

To help identify users and entities currently at risk in your organization, the **Entity Count Overview** widget displays the number of entities involved in risky behaviors, by entity type, along with their risk counts based on the last analytics run. When you click an entity type in the widget, the Entities page opens in the ArcSight Intelligence UI, where additional information for the selected entity type is displayed.

Overall Risk Level

Requires ArcSight Intelligence be deployed. This widget is not available in a SaaS environment.

To help understand the general risk in your organization, the **Overall Risk Level** widget displays the trending risk of the organization based on the last analytics run.

Productivity

Requires ArcSight ESM be deployed. This widget is not available in a SaaS environment.

To help managers optimize analyst activity for the specified time range, the **Productivity** widget incorporates several elements related to SOC productivity:

Case Closure Velocity
Shows the current rate of case closure per week based on the target velocity for all owners and owner groups. For example, you might expect teams to close at least 5 cases per week. The dotted line in the graph represents the target.

The trend indicates whether the velocity fails to meet or exceeds the target rate compared to the previous week. The velocity is based on when cases were created.

Highest Velocity
Represents the owner that currently has the fastest closure rate per week. You can also see the total number of cases assigned to the owner by severity.

The trend indicates whether the velocity fails to meet or exceeds the target rate compared to the previous week. The velocity is based on when cases were assigned to the owner.

Productivity by Owner Groups
Lists the owner groups that currently have the highest average number of cases closed per week. It also identifies which owner in the group has the highest velocity.

You can observe the average number of cases closed and whether the rate is trending up or down. The colored bar indicates the volume of cases by severity. By default, the widget displays data according to the specified time range.

Productivity (SOAR)

Requires data from ArcSight SOAR.

To help managers optimize analyst activity for the specified time range, the **Productivity (SOAR)** widget incorporates several elements related to SOC productivity. You can change the widget's properties to select an available option from the **Number of Groups** drop-down list:

Case Closure Velocity

Shows the current rate of case closure per week based on the target velocity for all owners and owner groups. For example, you might expect teams to close at least 5 cases per week. The dotted line in the graph represents the target.

The trend indicates whether the velocity fails to meet or exceeds the target rate compared to the previous week. The velocity is based on when cases were created.

Highest Velocity

Represents the owner that currently has the fastest closure rate per week. You can also see the total number of cases assigned to the owner by severity. The trend indicates whether the velocity fails to meet or exceeds the target rate compared to the previous week. The velocity is based on when cases were assigned to the owner.

Productivity by Owner Groups

Lists the owner groups that currently have the highest average number of cases closed per week. It also identifies which owner in the group has the highest velocity.

You can observe the average number of cases closed and whether the rate is trending up or down. The colored bar indicates the volume of cases by severity.

By default, the widget displays data according to the specified time range.

SOAR Average KPI for Event

Requires data from ArcSight SOAR and that ArcSight ESM be deployed. This widget is not available in a SaaS environment.

The **SOAR Average KPI For Event** widget provides the SOC Manager an overview for the volume of events in the specified time range that transition from initial analysis of events from source devices through correlation to case creation. The widget also shows the percentage of change between each state.

Correlated Event Count

Shows the number of alerts created from an ESM alert source that you must handle manually, without the use of ArcSight correlation.

Found

Indicates the reduction in the number of items that you must handle manually. This data includes the correlation events generated by rules that monitor events from source devices, as well as events generated by ArcSight components. For typical correlation rule configurations, the data usually represents a reduction in the number of items. However, the number of items might increase in the case of unusual configurations.

Created

Represents the number of cases created within a time range, based on correlation event activity, content, or systems detecting what is significant, and also manual assessments.

SOAR Case Breakdown - Severity

Requires data from ArcSight SOAR.

The **SOAR Case Breakdown - Severity** widget displays the number or percentage of cases by their **Severity**. The widget always shows data **As of Now**, regardless of the specified time range for the dashboard. By default, the widget shows data for total open cases. You can change the widget's properties to select or deselect a severity type. You can also create custom severities. The system, however, does not limit the number of custom severities that you can create.

SOAR Case Load

Requires data from ArcSight SOAR.

To help managers balance the amount of work assigned to case owners, the **SOAR Case Load** widget provides several case management metrics:

- Average number of cases each owner closes per week.
- Projection of the number of cases per severity that the owner might not be able to close, based on the configured target, the time elapsed since the cases were opened, and the average velocity of the owner. This assumes that owners work on cases in severity order, from highest to lowest.

> Estimation of the time required to close all cases is set in the **Severity Editor** when you configure case severities.

By default, the widget shows the data for total open, assigned cases for the top three members of the group based on their average number of cases per week. You can filter the data by specific Owner Groups. The metrics are based on the specified time range and the target number of cases that you expect the owners to close per Severity

For best use of this widget, it is recommended that you create one Case Load widget per owner group. In this way, you will see details for members of the owner group.

SOAR Case Status

Requires data from ArcSight SOAR.

The **SOAR Case Status** widget displays the number of cases by their Statuses. The widget always shows data **As of Now**, regardless of the specified time range for the dashboard.

By default, the widget shows data for All cases. You can however change the widget's properties to select or deselect one or more Status types.

SOAR Threat Analysis Funnel

Requires data from ArcSight SOAR and that ArcSight ESM be deployed. This widget is not available in a SaaS environment.

The **SOAR Threat Analysis Funnel** widget provides the SOC Manager an overview for the volume of events in the specified time range that transition from initial analysis of events from source devices through correlation to case creation. The widget also shows the percentage of change between each state.

Analyzed
Shows the number of **events**, from source devices, that you must handle with the use of ArcSight correlation.

Found
Indicates the reduction in the number of items that you must handle manually. This data includes the **correlation events** generated by rules that monitor events from source device as well as events created by ArcSight components. For typical correlation rule configurations, the data usually represents a reduction in the number of items. However, the number of items might increase in the case of unusual configurations.

> Base and Correlation Event counts and Created Case counts come from the *ArcSight ESM* API and *ArcSight SOAR* respectively.

SOAR Case Timeline

Requires data from ArcSight SOAR.

The **SOAR Case Timeline** widget shows changes in the volume of cases over a specified time range. By default, the widget filters the data according to the Severity category assigned in SOAR. However, you can also choose to view trends for other case states, such as cases closed by assigned or unassigned sub-filters.

To observe the breakdown of cases associated with a specific date, you can hover over any location within the timeline. You can also zoom in to view a particular time range, either using the magnifier icons or by clicking and dragging within the graph.

SOAR Top Playbooks Executed

Requires data from ArcSight SOAR.

The **SOAR Top Playbooks Executed** widget displays the execution count of the playbooks over alerts created.

By default, the widget shows data for top 5 playbooks. The widget helps managers understand the count of playbooks executed over each alert.

You can view the number of playbooks for a given date period and Top N Playbooks such as top 5, top 10.

SOAR Trend - Mean Time To Resolve

Requires data from ArcSight SOAR.

The **SOAR Trend - Mean Time to Resolve** widget displays the amount of average time it took to resolve a malicious attack.

You can change the widget's properties to view different classifications of attacks and their statuses.

To observe the breakdown of cases associated with a specific date, you can hover over any location within the timeline. You can also zoom in to view a particular time range, either using the magnifier icons or by clicking and dragging within the graph.

SOAR Trend - Mean Time To Response

Requires data from ArcSight SOAR.

The **SOAR Trend - Mean Time to Response** widget displays the amount of average time it took to respond to a malicious attack.

You can change the widget's properties to view different classification of attacks and their statuses.

To observe the breakdown of cases associated with a specific date, you can hover over any location within the timeline. You can also zoom in to view a particular time range, either using the magnifier icons or by clicking and dragging within the graph.

SOAR Trend - Playbooks Executed

Requires data from ArcSight SOAR.

The **SOAR Trend - Playbooks Executed** widget displays the number of times a playbook is executed by its execution date.

By default, the widget shows data for 5 playbooks. The widget helps managers understand the number of playbooks executed everyday.

You can change the widget's properties to view number of playbooks for a given date period.

To observe the breakdown of playbooks associated with a specific date, you can hover over any location within the timeline. You can also zoom in to view a particular time range, either using the magnifier icons or by clicking and dragging within the graph.

Threat Analysis Funnel

Requires ArcSight ESM be deployed. This widget is not available in a SaaS environment.

The **Threat Analysis Funnel** provides the SOC Manager an overview for the volume of events in the specified time range that transition from initial analysis of events from source devices through correlation to case creation. The widget also shows the percentage of change between each state.

Analyzed
Shows the number of **events**, from source devices, that would need to be handled manually without the use of ArcSight correlation.

Found
Indicates the reduction in the number of items that you would need to handle manually. This data includes the **correlation events** generated by rules that monitor events from source device as well as events created by ArcSight components. For typical correlation rule

configurations, the data usually represents a reduction in the number of items. However, the number of items might increase in the case of unusual configurations.

Created

Represents the number of **cases** created within the time range, based on correlation event activity, content or systems detecting what's significant, and manual assessments.

Top Risky Entities

Requires ArcSight Intelligence be deployed. This widget is not available in a SaaS environment.

To help identify the riskiest entities in your organization, the **Top Risky Entities** widget provides a list of the top risky entities, by entity type, based on the last analytics run. By default, the widget displays the top 5 risky users. If you need to view the top risky entities for another entity type, then, as part of this widget's properties, you can change the filter to select the entity type and the number of entities you want displayed in the list. When you click an entity in the widget, the Explore page opens in the ArcSight Intelligence UI, with the selected entity's name applied to the **anomalies and violations** filter.

Accessing ArcMC

Available only with ArcSight capabilities. Not available in a SaaS environment.

Select **ARCMC**.

ArcSight Management Center (ArcMC) enables you to manage and monitor ArcSight infrastructure components, particularly useful when you have a large deployment ArcSight connectors. From the **ArcMC dashboards**, you can view the health and status of the components that ArcMC manages. The **Bulk Operations** feature allows you to modify the properties of, ensure the security of, gather log information about, and restart managed components.

Accessing Bulk Operations

You must have the ArcSight Management Center deployed to use Bulk Operations.

Select **ARCMC** > **Bulk Operations**.

Bulk Operations enables you to view and manage collectors, hosts and locations of hosts, and Transformation Hubs. You can modify the properties of, ensure the security of, gather log information about, and restart managed components.

Accessing ArcMC Dashboards

Select **ARCMC** > **Dashboards**.

The dashboards enable you to view the health and status of the components that ArcMC manages.

Managing Your Profile

Select [your_ID] > My Profile.

You can manage your account settings and review your assigned roles, permissions, and groups. Also, configure your preferred default settings for product behavior and interface theme.

Manage Your Account

Select **[your_ID]** > **My Profile** > **MY PROFILE.**

You can change your account settings. However, you cannot change your password in Fusion if your enterprise uses an external authentication method.

Configure Your User Preferences

Select **[your_ID]** > **My Profile** > **PREFERENCES**.

Some deployed capabilities enable you to configure preferences for commonly used settings. For example, if you regularly use the same fieldset for a Search, you can specify that set as your preferred default.

Configure Search Preferences

To reduce the time required to create and manage searches, configure Search to use your preferred settings. You can always override your preferences as needed when you create a search.

Default Fieldset

Specifies the fieldset that you regularly use for a search. The default value is *Base Event Fields*.

Default View

Specifies whether you want the Events Table to display results in the **Grid View** or **Raw View**. The default value is *Grid View*.

Time Zone

Instructs Search to adjust the timestamp for events to the chosen time zone:

- Browser
- Database
- Custom

To specify the type of timestamp that you want to use, modify the preference for **Base Searches On**.

Date / Time Format

Specifies the format of dates and times that you want Search to use. The default is *YYYY/MM/DD*.

For example, you might want to use the same format that you have already configured for your browser. Alternatively, you might prefer a format like *MM/DD/YYYY HH:MM:SS*.

Default Time Setting

Specifies the time range within which you want Search to find events. The default is *Last 30 minutes*.

- Dynamic

  If you prefer to use a dynamic time range, you must also specify the **Start** and **End** times. For example, specify **$Now - 30m** and **$Now** respectively.

- Static

  If you use different time settings for each search that you create, you might want to select this option for your preference. The default is the preset value of *Last 30 minutes*.

- Preset

  If you prefer to use a preset time range, you must also specify a preset value. For example, **Last 24 hours**.

Base Searches On

Specifies the timestamp associated with the events that you want to find:

- Normalized Event Time

- Device Receipt Time

- Database Receipt Time

Default is *Normalized Event Time*.

Search Expires in

Specifies how often you want saved searches to expire, and thus for the system to remove them from the system. You can specify a value between 1 and 365. Your System Admin might also specify a different range. To specify a value above 365, check with your System Admin. The expiration date resets whenever you access the search. Resetting the expiration date includes resuming or re-running the search, as well as saving the search and changing its settings.

The default value is *7 days*. Alternatively, if you have the *Never Expire Search Results* permission, you can choose for a search to never expire. When you create or edit a search, you can override this default setting.

Session Search Expires In

Specifies how often you want searches to expire, and thus be removed from the system. You can specify a value between 1 and 120. Your System Admin may also specify a different range. to specify a value above 120, check with your System Admin. You can also choose to never remove a search. The expiration time resets whenever you change or run the search.

When you create or edit a search, you can override this default setting.
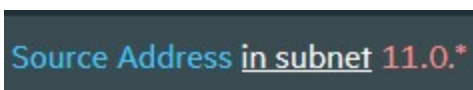
Maximum Results for a Search

Specifies the maximum number of events that the Search will return. You can specify a value between 1000 and 10 million. The default is *300,000*. When creating a search, you can override this preference.

Your admin can configure a system-level setting that controls the maximum number of searches (with a limit of 10 million) for all instances of Fusion. If you enter a value outside of the system-level setting, you will receive an error message indicating that your preferred default cannot exceed the system setting. For information about setting a global search limit in a non-SaaS environment, see Configuring the Deployed Capabilities in the *Administrator's Guide for the ArcSight Platform*. If you are a SaaS customer, reach out to Support to increase the search limit.

Highlight Query Syntax

Specifies whether you want Search to use color to differentiate the syntax terms from the operators and functions within the query. For example, in the figure below, Search displays the variable *Source Address* in blue, the value *11.0.** in red, and the operator *in subnet* in white.

Review Your Roles and Permissions

Select [your_ID] > My Profile > ROLES & PERMISSIONS.

You can review the roles assigned to your account, and the permissions associated with each role.

Review Your Group Assignments

Select **[your_ID]** > **My Profile** > **GROUPS**.

You can review the account groups to which you belong, as well as the manager of the group.

Set Your Default Theme

Select **[your_ID]** > **My Profile** > **THEMES**.

You can specify which theme you want to use as a default for the interface. The **dark** and **light** themes are built-in. However, an administrator can create additional themes from which you can choose. If an administrator hides or deletes your default theme, Fusion automatically changes your default theme to the built-in *Dark* theme.

Choose the theme that you want to use, then select **APPLY THEME**.

VII Appendices

The appendices in this guide provide additional information or guidance for using the features and functions for this product.

Agent Fields

Substitute the following labels in the agent category:

For the field that you want to add...	You should choose...
Agent Address	agentAddressBin
Agent DNS Domain	agentDnsDomain
Agent Hostname	agentHostName
Agent ID	agentId
Agent Mac Address	agentMacAddressBin
Agent NT Domain	agentNtDomain
Agent Receipt Time	agentReceiptTime
Agent Severity	agentSeverity
Agent Timezone	agentTimeZone
Agent Translated Address	agentTranslatedAddressBin
Agent Translated Zone External ID	agentTranslatedZoneExternalID
Agent Translated Zone URI	agentTranslatedZoneURI
Agent Type	agentType
Agent Version	agentVersion
Agent Zone External ID	agentZoneExternalID
Agent Zone URI	agentZoneURI

Category Fields

Substitute the following labels in the category:

Category Behavior	categoryBehavior
Category Device Group	categoryDeviceGroup
Category Device Type	categoryDeviceType
Category Object	categoryObject
Category Outcome	categoryOutcome
Category Significance	categorySignificance
Category Technique	categoryTechnique
Version	version

Correlation Fields

Substitute the following labels in the `correlation` category:

Substitute the following labels in the correlation category:	You should choose...
Base Event Ids	correlated_event_id
Correlated Event Id	generatorURI
Generator External ID	generatorExternalID
Generator URI	base_event_ids
Priority	priority

Destination Fields

Substitute the following labels in the `destination` category:

For the field that you want to add...	You should choose...
Destination Address	destinationAddressBin
Destination DNS Domain	destinationDnsDomain
Destination Geo Country Code	destinationGeoCountryCod
Destination Geo Latitude	destinationGeoLatitude
Destination Geo Longitude	destinationGeoLongitude
Destination Geo Postal Code	destinationGeoPostalCode
Destination Geo Region Code	destinationGeoRegionCode
Destination Geolocation Info	destinationGeoLocationInfo
Destination Hostname	destinationHostName
Destination Mac Address	destinationMacAddressBin
Destination NT Domain	destinationNtDomain
Destination Port	destinationPort
Destination Process ID	destinationProcessId
Destination Process Name	destinationProcessName
Destination Service Name	destinationServiceName
Destination Translated Address	destinationTranslatedAddressBin
Destination Translated Port	destinationTranslatedPort
Destination Translated Zone External ID	destinationTranslatedZoneExternalID
Destination Translated Zone URI	destinationTranslatedZoneURI
Destination User ID	destinationUserId
Destination User Privileges	destinationUser Privileges
Destination Username	destinationUserName
Destination Zone External ID	destinationZoneExternalID
Destination Zone URI	destinationZoneURI

Device Fields

Substitute the following labels in the device category:

For the field that you want to add...	You should choose...
Device Action	deviceAction
Device Event Class ID	deviceEventClassId
Device Address	deviceAddressBin
Device Asset ID	deviceAssetID
Device Direction	deviceDirection
Device DNS Domain	deviceDnsDomain
Device Domain	deviceDomain
Device Event Category	deviceEventCategory
Device Inbound Interface	deviceInboundInterface
Device Event Class ID	deviceEventClassId
Device External ID	deviceExternalId
Device Facility Hostname	deviceFacility
Device Hostname	deviceHostName
Device Mac Address	deviceMacAddressBin
Device NT Domain	deviceNtDomain
Device Outbound Interface	deviceOutboundInterface
Device Process ID	deviceProcessId
Device Process Name	deviceProcessName
Device Product	deviceProduct
Device Receipt Time	deviceReceiptTime
Device Severity	deviceSeverity
Device Timezone	deviceTimeZone
Device Translated Address	deviceTranslatedAddressBin
Device Translated Zone External ID	deviceTranslatedZoneExternalID
Device Translated Zone URI	deviceTranslatedZoneURI
Device Version	deviceVendor

Device Version	deviceVersion
Device Zone External ID	deviceZoneExternalID
Device Zone URI	deviceZoneURI
Normalized Event Time	normalizedEventTime

Device Custom Fields

Substitute the following labels in the `device custom` category:

For the field that you want to add...	You should choose...
Device Custom Date 1	deviceCustomDate1
Device Custom Date 1 Label	deviceCustomDate1Label
Device Custom Date 2	deviceCustomDate2
Device Custom Date 2 Label	deviceCustomDate2Label
Device Custom Descriptor ID	deviceCustomDescriptorId
Device Custom Floating Point 1	deviceCustomFloatingPoint1
Device Custom Floating Point 1 Label	deviceCustomFloatingPoint1Label
Device Custom Floating Point 2	deviceCustomFloatingPoint2
Device Custom Floating Point 2 Label	deviceCustomFloatingPoint2Label
Device Custom Floating Point 3	deviceCustomFloatingPoint3
Device Custom Floating Point 3 Label	deviceCustomFloatingPoint3Label
Device Custom Floating Point 4	deviceCustomFloatingPoint4
Device Custom Floating Point 4 Label	deviceCustomFloatingPoint4Label
Device Custom Number 1	deviceCustomNumber1
Device Custom Number 1 Label	deviceCustomNumber1Label
Device Custom Number 2	deviceCustomNumber2
Device Custom Number 2 Label	deviceCustomNumber2Label
Device Custom Number 3	deviceCustomNumber3
Device Custom Number 3 Label	deviceCustomNumber3Label
Device Custom String 1	deviceCustomString1
Device Custom String 1 Label	deviceCustomString1Label
Device Custom String 2	deviceCustomString2
Device Custom String 2 Label	deviceCustomString2Label
Device Custom String 3	deviceCustomString3
Device Custom String 3 Label	deviceCustomString3Label
Device Custom String 4	deviceCustomString4

For the field that you want to add...	You should choose...
Device Custom String 4 Label	deviceCustomString4Label
Device Custom String 5	deviceCustomString5
Device Custom String 5 Label	deviceCustomString5Label
Device Custom String 6	deviceCustomString6
Device Custom String 16 Label	deviceCustomString6Label
Device CustomIPv6 Address 1	deviceCustomIPv6Address1Bin
Device CustomIPv6 Address 1 Label	deviceCustomIPv6Address1Label
Device CustomIPv6 Address 2	deviceCustomIPv6Address2Bin
Device CustomIPv6 Address 2 Label	deviceCustomIPv6Address2Label
Device CustomIPv6 Address 3	deviceCustomIPv6Address3Bin
Device CustomIPv6 Address 3 Label	deviceCustomIPv6Address3Label
Device CustomIPv6 Address 4	deviceCustomIPv6Address4Bin
Device CustomIPv6 Address 4 Label	deviceCustomIPv6Address4Label

Event Fields

Substitute the following labels in the event category:

For the field that you want to add...	You should choose...
Application Protocol	applicationProtocol
Base Event Count	baseEventCount
Bytes In	bytesIn
Bytes Out	bytesOut
Crypto Signature	cryptoSignature
Customer External ID	customeExternalID
Customer URI	customerURI
End Time	endTime
Event ID	eventId
Event Outcome	eventOutcome
External Id	externalID
Locality	locality
Message	message
Name	name
Originator	originator
Reason	reason
Start Time	startTime
Transport Protocol	transportProtocol
Type	type

Extension Fields

Substitute the following labels in the `extension` category:

For the field that you want to add...	You should choose...
Extra Fields	extraFields
Storage Group	storageGroup

File Fields

Substitute the following labels in the `file` category:

For the field that you want to add...	You should choose...
File Create Time	fileCreateTime
File Hash	fileHash
File ID	fileId
File Modification Time	fileModificationTime
File Name	fileName
File Path	filePath
File Permission	filePermission
File Size	fileSize
File Type	fileType

Flex Fields

Substitute the following labels in the `flex` category:

For the field that you want to add...	You should choose...
Flex Date 1	flexDate1
Flex Date 1 Label	flexDate1Label
Flex Number 1	flexNumber1
Flex Number 1 Label	flexNumber1Label
Flex Number 2	flexNumber2
Flex Number 2 Label	flexNumber2Label
Flex String 1	flexString1
Flex String 1 Label	flexString1Label
Flex String 2	flexString2
Flex String 2 Label	flexString2Label

OldField Fields

Substitute the following labels in the `oldfield` category:

For the field that you want to add...	You should choose...
Old File Create Time	oldFileCreateTime

Old File Fields

Substitute the following labels in the `old file` category:

For the field that you want to add...	You should choose...
Old File Hash	oldFileHash
Old File ID	oldFileId
Old File Modification Time	oldFileModificationTime
Old File Name	oldFileName
Old File Path	oldFilePath
Old File Permission	oldFilePermission
Old File Size	oldFileSize
Old File Type	oldFileType

Request Fields

Substitute the following labels in the `request` category:

For the field that you want to add...	You should choose...
Request Client Application	requestClientApplication
Request Context	requestContext
Request Cookies	requestCookies
Request Method	requestMethod
Request URL	requestUrl
Request URL FileName	requestUrlFileName
Request URL Query	requestUrlQuery

Source Fields

Substitute the following labels in the `source` category:

For the field that you want to add...	You should choose...
Source Address	sourceAddressBin
Source DNS Domain	sourceDnsDomain
Source Geo Country Code	sourceGeoCountryCode
Source Geo Latitude	sourceGeoLatitude
Source Geo Longitude	sourceGeoLongitude
Source Geo Postal Code	sourceGeoPostalCode
Source Geo Region Code	sourceGeoRegionCode
Source Geolocation Info	sourceGeoLocationinfo
Source Hostname	sourceHostName
Source Mac Address	sourceMacAddressBin
Source NT Domain	sourceNtDomain
Source Port	sourcePort
Source Process ID	sourceProcessId
Source Process Name	sourceProcessName
Source Service Name	sourceServiceName
Source Translated Address	sourceTranslatedAddressBin
Source Translated Port	sourceTranslatedPort
Source Translated Zone External ID	sourceTranslatedZoneExternalID
Source Translated Zone URI	sourceTranslatedZoneURI
Source User ID	sourceUserId
Source User Privileges	sourceUser Privileges
Source Username	sourceUserName
Source Zone External ID	sourceZoneExternalID
Source Zone URI	sourceZoneURI

Publication Status

Released: January 5, 2023

Updated: Wednesday, April 26, 2023

Support

Contact Information

Phone	A list of phone numbers is available on the Technical Support Page: https://softwaresupport.softwaregrp.com/support-contact-information
Support Web Site	https://softwaresupport.softwaregrp.com/
ArcSight Product Documentation	https://www.microfocus.com/documentation/arcsight/

Send Documentation Feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this computer, click the link above and an email window opens with the following information in the subject line:

Feedback on Fusion in the ArcSight Platform User's Guide (Fusion in the ArcSight Platform User's Guide 1.6.0)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to Documentation-Feedback@microfocus.com.

We appreciate your feedback!