Hewlett Packard Enterprise

HPE Security ArcSight Investigate

Software Version: 1.01

User's Guide

June 15, 2017

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

The network information used in the examples in this document (including IP addresses and hostnames) is for illustration purposes only.

HPE Security ArcSight products are highly flexible and function as you configure them. The accessibility, integrity, and confidentiality of your data is your responsibility. Implement a comprehensive security strategy and follow good security practices.

This document is confidential.

Restricted Rights Legend

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2017 Hewlett Packard Enterprise Development, LP

Follow this link to see a complete statement of copyrights and acknowledgements:
https://community.saas.hpe.com/t5/Discussions/Third-Party-Copyright-Notices-and-License-Terms/td-p/1589228

Support

Contact Information

Phone	A list of phone numbers is available on the HPE Security ArcSight Technical Support Page: https://softwaresupport.hpe.com/documents/10180/14684/esp-support-contact-list
Support Web Site	https://softwaresupport.hpe.com
Protect 724 Community	https://community.saas.hpe.com/t5/ArcSight/ct-p/arcsight

Contents

Chapter 1: Introduction 1
- How ArcSight Investigate works 1
- Features and benefits 4
- Workflow summary 5
- Editing my profile 9

Chapter 2: User management 10
- Managing users 10
  - Creating a user 10
  - Importing users from ESM 11
  - Managing user groups 13
    - What is a user group? 13
    - Creating a user group 14
    - Adding users to a group 14
    - Removing a user from a group 15
    - Adding or removing a group manager 16
    - Deleting a group 16
  - Searching for a user 16
  - Managing a user's account 17
- Managing roles 18
  - What is a role? 18
  - Creating a role 18
  - Editing a role 19
  - Deleting a role 20
  - Assigning a role to a user 20
  - Removing a role from a user 21
  - Built-in roles 22
  - Permissions 22

Chapter 3: Processes 24
- Terminating processes 24

Chapter 4: Managing dashboard widgets 25
- Adding a widget to the Dashboard 25
- Deleting a widget from the Dashboard 26

Chapter 5: Searching event data 27
- Searching events 27
- Searching events in ESM 31

Managing search-results fieldsets 35
Creating a fieldset for search results 35
Editing a fieldset for search results 36
Deleting a fieldset for search results 37
Viewing search results for a time range 38
Charting search-results data 38
Creating line, bar, column, and area charts 39
Creating a pie chart 44
Creating a scatter plot chart 46
Editing search-results charts 48
Deleting search-result charts 50
Managing search results information 51
Viewing the most and least common values for an event record field 51
Pinning field columns to help analyze events 52
Viewing all fields of an event 52
Viewing select event data 53
Exporting search results 54
Appendix A: FAQs 56
Can I pin a field column in order to compare it against other field values? 56
Can I export search-results data to an Excel file? 56
How much search-result data can I view? 56
Can I view the most and least common values for a search-results field? 57
Can I use SQL to specify query input? 57
Can I use a SIEM with ArcSight Investigate? 57
Can I apply User Behavior Analytics to the Hadoop data lake used by ArcSight Investigate? 57

Send Documentation Feedback 59

Glossary 60

Chapter 1: Introduction

ArcSight Investigate enables you to search, analyze, and visualize machine-generated data gathered from such entities as websites, applications, sensors, and devices that comprise your IT infrastructure or business.

After Investigate ingests the data stream of individual events from Event Broker, you can view and search.

You can use the English-like search language to create searches.

How ArcSight Investigate works

ArcSight Investigate is a high-capacity, threat-investigation solution that enables you to search through and analyze (not just search) vast amounts of event data for anomalies associated with such entities as users, IP addresses, and network assets. Information yielded from a search can help you detect and investigate breaches before substantial damage can be done to your organization. From this, you can also discover contextual information regarding these entities and the effectiveness of security policies and rules, and security applications.

ArcSight Event Broker and ArcSight SmartConnectors are essential parts of the Investigate solution. Connectors send normalized and categorized CEF events to the ArcSight Event Broker topic (`eb-cef`). Events are transformed to Apache Avro format and then consumed by the Vertica Kafka scheduler and then loaded to the HPE Vertica database.

The Vertica scheduler pulls events from a topic and then loads the events into the Vertica database. Investigate reads the events from the Vertica database and then displays them in the Search page.

Investigate can extend the ArcSight Enterprise Security Manager (ESM) application in order to further investigate events in an active channel. ESM generates a URL that opens Investigate, with query input based on the data selected in the active channel.

The Search function makes it possible for you to investigate security incidents. It uses a microservices-based architecture, where it is possible to isolate different components (microservices) using docker container technology. This technology enables you to install and configure complex software and package it in an isolated environment—a container, then deploy these containers in any environment making it possible to maintain multiple servers even at a customer-controlled host or the cloud. Even though the microservices are isolated in containers, they still need to interact with each other in order for the application to run as a whole. This interaction is usually done by REST calls between the different microservices.

The Search function is composed of four basic components:

- Search UI

  The Search page is where you start an investigation. It is composed of the **Search** field, Filter field, Timeline, data visualization charts, and Events table:

- Search backend

  The Search backend saves searches, user preferences, and proxy search requests to Search engine. The REST API is used to implement this.

- Search engine

  Search engine is a scalable server-side application that is responsible for executing and caching large search queries in the Vertica database.

- Vertica database

  The database serves as the main data store, as well as a cache.

Investigate pages

- Dashboard — Where you view data visualization charts and text boxes for note taking.
- Search — Where you perform searches on events and manage this process.
- Admin — Where you set up users and establish user rights.

Roles and functions

- Security analyst

  Functions

  - Search events generated in your network (see "Searching events" on page 27).
  - Send an event from an ArcSight Enterprise Security Manager (ESM) channel to Investigate for further investigation (see "Searching events in ESM" on page 31).
  - Manage search-results fieldsets (see "Managing search-results fieldsets" on page 35)
  - Visualize search result data (see "Charting search-results data" on page 38).
  - Manage the display of search results information to better detect and analyze anomalies (see "Managing search results information" on page 51).
  - Manage dashboard widgets (see "Managing dashboard widgets" on page 25).

  Product access

  Investigate Console (Dashboard page and Search page)

- Security engineer

  Functions

  - Can operate as a security analyst.

  Product access

  Investigate Portal (Admin page) and Investigate Console

- System admin

Functions

- Can operate as a security analyst.
- Installs and deploys Investigate.
- Can add and remove analysts and architects.

Product access

Search, Dashboard, and Admin pages

Features and benefits

The following are the main features and benefits of ArcSight Investigate.

- **Event scalability**

  Investigate can collect and search on multiple data, including security, IT, OT and IoT. Investigate can receive data from all ArcSight SmartConnectors. Using Event Broker, the application can effectively distribute event ingestion across multiple machines. Receiving heterogeneous data and nearly linear scaling make it quick and easy to expand your security posture.

- **Search**

  Search is the primary way to navigate data in Investigate. The search is contextual and has auto-suggest capability to help you specify the query input. Therefore, there is no need to learn a complex query language or schema. This can boost the productivity of analysts. Using Vertica, information retrieval is extremely rapid, making a search up to 10x faster than competing products. You can write a search to retrieve events from an index, use statistical commands to calculate metrics and generate charts, search for specific conditions within a rolling time window, identify patterns in your data, predict future trends, and so on. Data visualization charts can be added to a search in order to better understand search-results data.

- **Indexing**

  Investigate indexes machine data. This includes data streaming from packaged and custom applications, application servers, web servers, databases, networks, virtual machines, telecoms equipment, operating systems, sensors, and so on, that make up your IT infrastructure. The maximum indexing volume depends on the Investigate license.

- **Data Analysis**

  Investigate enables you to conduct a security investigation by filtering, comparing, visualizing, and analyzing event data dynamically. You are able to expedite the investigation process with quick and easy data analysis, deriving insights without any complexity. Investigate provides precise investigation outcomes through pre-defined queries (and fieldsets) for security use cases; therefore, improving SOC efficiency and reducing threat posture.

- **Charting**

You can graphically represent search-results data using the chart editor. This editor enables you to map attributes defined by data-model objects to a chart data visualization without having to write the searches to generate them. A chart is saved with a search and can be added to dashboards.

- **Dashboard**

  The dashboard can be comprised of search-results charts or text panels, or a combination of these. The dashboard can display the results of completed searches as well as data from real-time searches that run in the background.

Workflow summary

As a security analyst, you need to search event data as part of an ad-hoc analysis and investigation of security incidents and threats. Below, is a summary of this workflow and how ArcSight Investigate can help you in this effort.

Typically, the workflow originates from the Search page. However, there could be alternative starting points in the workflow, such as the Dashboard page.

Set up Investigate.

Perform the installation with a root user or with a user that has pseudo rights.

1. Copy the installation package to the master server (see "Install ArcSight Investigate and Event Broker using the ArcSight Installer application" in the *HPE Security ArcSight Investigate Deployment Guide*).
2. Configure Vertica, SMTP, and general settings (see "Configure ArcSight Investigate components" in the *HPE Security ArcSight Investigate Deployment Guide*).

Open Investigate.

When you log in for the first time, you need to create the first user in the system. This user is assigned with the system admin role.

3. Launch Investigate from an approved browser using the URL (see "Supported browsers" in the *HPE Security ArcSight Investigate Deployment Guide*):

   `https://<master_node_IP>`
4. From the Welcome page, specify information about yourself.
5. Create a system admin (see "Creating a user" on page 10).
6. From the Login page, specify your credentials.

   From the welcome email that you received, a link was provided for you to create a password.

Add security analysts.

Location: Left navigation > Admin > User Groups

7. Click **Create User** and then specify the email, name, group(s), and roles for the security analyst (see "Creating a user" in the *ArcSight Investigate Administrator's Guide*).

   Security analysts are identified by their email. You can only create one analyst with a specific email. By default, all analysts are assigned to the All Users group.

Create a search.

Some commonly used search queries are available as "canned" queries (see "Searching events" on page 27).

Location: Left navigation > Search button > Search page

8. Click **New Search** and then accept the default search name or rename it (see "Searching events" on page 27).

   - Investigate automatically assigns a name to a new search in the format of `Search x + 1`.
   - To save your work throughout the search procedure, click **Save** from the task bar.

Specify the query input.

The context of a search is twofold: (1) Investigate an alert or incident and (2) gather and compile data for analysis (see "Searching event data" on page 27).

Location: Left navigation > Search button > Search page

9. From the **Search** field, specify the desired query input (see "Searching events" on page 27).

   - The query can either be full text, natural language, or contextual.
   - An event search consists of specifying query input, search-result fields, and the time period to search within.
   - You can filter search results for more specific information.
   - To use "canned" queries, type # and then select the desired query.

Make search-result fields available.

The default fieldset contains the 52 most common event fields. These fields are available for creating data visualization charts and for viewing in the Events table. Each field (column) in this table can provide the 10 most and least common values.

Location: Left navigation > Search button > Search page

10. To create a fieldset, click the fieldset button and then specify the desired fields (see "Creating a fieldset for search results" on page 35).
11. To edit a fieldset, click the fieldset button, select **Edit this set** from the drop-down, and then specify the desired fields (see "Editing a fieldset for search results " on page 36).
12. To delete a fieldset, click the fieldset button, select **Edit this set** from the drop-down, and then click

Delete (see "Deleting a fieldset for search results " on page 37).

- If only a single search is using the fieldset, then the fieldset will be deleted and the default fieldset used in its place.
- If two or more searches are using the fieldset, then the fieldset cannot be deleted.

Specify the search time period.

By default, the time range is for the last 30 minutes (current time minus 30 minutes).

Location: Left navigation > Search button > Search page

13. From the time drop-down, specify the time period to be searched (see "Searching events" on page 27).

    The Quick Ranges feature provides a convenient way to select common search periods ranging from the past minute up to a year ago—or the whole time range where events occurred.

14. Click **Search**.

    To cancel the search, click **X** in the **Search** field.

To narrow your search, do any necessary filtering.

Filtering search results can be accomplished by creating a filter, specifying a time segment in the Timeline, and by adjusting the time range to be searched.

Location: Left navigation > Search button > Search page

15. Click **Filter** and then make the appropriate selections from the **Select** drop-downs (see "Searching events" on page 27).

    To add another filter, click **Add Filter**.

    Other filtering methods:

    - Drag and drop a statement from the Events table to the "Filter" area.

      Upon drop, the "Filter" area hightlights if the statement is acceptable.

    - Right click a data cell in the Events table and then choose **Use As A Filter**.

    For both cases, Investigate adds a new filter (row) to the "Filter" area.

16. From the "Timeline" area, ensure that the range selector is on and then specify the desired time range ("Viewing search results for a time range" on page 38).

    - This filter focuses on the search results for a specific time range within the specified search time period.
    - The new results are reflected in any data visualization and the Events table.

To better understand search-result data, represent it graphically.

You can create up to 10 data visualization charts for a search.

Location: Left navigation > Search button > Search page > Visualize

17. To chart search-result data, expand the "Visualize" area and then select the desired chart type (see Visualizing search-results data).

    Using the ..., you can add the chart to the Dashboard and rename the chart.

18. To enlarge the chart, click the chart.
19. To change the contents of a chart, see "Editing search-results charts" on page 48.
20. To delete a chart, see "Deleting search-result charts" on page 50.

To view select search results, organize data in the Events table.

Investigate lists search-result events in the Events table. The various fields of the event records are represented by the table column headers. Field columns are determined by the fieldset being used and the Show Columns feature.

Location: Left navigation > Search button > Search page > Events area

21. To view the most and least common values for an event field, right click and specify to view these values from the desired field column header (see "Viewing the most and least common values for an event record field" on page 51).

    The most and least common values for an event record field translates into the count and percentage of hits for the field value.

22. To help compare the column values against those of other columns, make the desired column sticky (see "Pinning field columns to help analyze events" on page 52).
23. To view all the fields of an event, click the arrow of the desired agent address (see "Viewing all fields of an event" on page 52).

    This feature enables you to quickly view the details of a single event without having to add all the fields of the fieldset.

24. To isolate an event for viewing, click the desired event row and specify view only (see "Viewing select event data" on page 53).

    Use the shift and control keys to select multiple events.

25. To limit the display of field columns, click and then deselect undesired fields.
26. To reorder field columns, click and drag desired columns to new positions.
27. To sort values for a field column, click the appropriate arrow in the desired column heading.

A single click in the column heading sorts values in ascending order while a double click sorts values in descending order.

28. To view events with a common field value, right click on the desired field value and search for the field value.

    Below the selected event record, Investigate list all the event records with a matching field value.

Editing my profile

About

Your profile includes the following:

- Personal details (name and email).
- Your roles and permissions.
- Fields to which you have access.
- Groups to which you belong.

Procedure

Location: Left navigation > [your name] > My Profile

- Make desired changes and then click **Save**.

Chapter 2: User management

ArcSight Investigate employs role-based access control, a method for regulating an individual user's access according to his or her role. This chapter includes information on managing users and roles in ArcSight Investigate.

This guide includes the following chapters:

"Managing users" below

"Managing roles" on page 18

Managing users

Typically, users are managed in groups. When a user is created in ArcSight Investigate, the user is automatically added to the All Users group (built-in group). You do not have to create additional groups in order to manage users.

The main user management tasks include:

- Creating a user:
  - Individually, as described in "Creating a user" below
  - Importing multiple users, as described in "Importing users from ESM" on the next page
- "Adding users to a group" on page 14
- "Assigning a role to a user" on page 20
- "Managing a user's account" on page 17

Creating a user

About

Required permissions:

- Create Users

Users are identified by their email. You can only create one user with a specific email.

By default, all users are assigned to the **All Users** group. Users cannot be removed from this group.

Procedure

Location: Left navigation > Admin > User Groups

1. In the top navigation pane, click **Create User**.
2. In the **Create User** dialog box, enter the following user details:

- Email
- First name
- Last name

3. In the **Groups** section, select the groups to which you want to add this user.

   Note: We recommend that you assign the user to a group that you manage, otherwise, you won't be able to manage the user.

4. In the **Roles** section, select the roles that you want to assign to this user.

   You can only assign roles that you have yourself. All other roles are displayed as read-only.

   Note: If you are creating an Admin user, in addition to the Admin role, make sure to assign the user roles that he can assign to the users that he creates. For example, the Guest and the User roles.

5. Click **Save**.

   The new user receives a welcome email that includes a link for creating a password for ArcSight Investigate. The link is valid for 24 hours.

6. To create another user, click **Save and Add Another**.

   The groups and roles remain selected for the next user that you create.

Importing users from ESM

About

You must have the System Admin role to import users.

You can import users from ESM using the ArcSight Investigate user interface.

Users are imported along with their groups, as defined in ESM.

The following fields are imported:

- First name
- Last name
- email
- Group

Users and groups are imported from the Users resource.

Import logic

- A user is identified by his or her email.
- Only users that adhere to the following conditions are imported:

  - Must have a valid email address.
  - Must be either a Web User or Normal User (user type).
  - Must be log in enabled.
- If a user does not have a first or last name in ESM, then the name is extracted from the user's email address. For example: if the user's email is *johndoe*@hpe.com, then the first and last name will both be *johndoe*. You can edit the user's name in ESM before you import or after you import, in ArcSight Investigate.
- All groups are imported, including deprecated groups. Deprecated groups are marked as deprecated in the import report.
- You can re-import from ESM as many times as you want. If a user has already been imported, his or her details will not be altered on re-import. However, if the user belongs to an additional group, then the user will be added to that group.
- If there are two groups with the same name, the name of one of the groups is concatenated with the parent group name.
- Empty groups are not created. If all the users in the group are invalid, the group is not created.

Once the import is complete, a report is displayed with all the users that were imported or that could not be imported and the groups that were created. You can click on the user's name to open the user's personal details. Click the Back button in the browser to return to the report.

Prerequisites

- ESM host name or IP address (as specified in the ESM Manager certificate)
- Port
- ESM credentials: Admin User ID and Password

Import is supported for ESM version 6.11 and higher.

If you have more than one ESM, perform the following procedure for each ESM server.

Procedure

Location: Left navigation > Admin > User Groups

1. From the top navigation pane, click **Import Users**.
2. Enter the following information:
   - ESM host name
   - Port (default 8443)
   - Admin User ID
   - Password
3. Select the roles that you want to assign to the imported users.

   You can edit roles for specific users after they are imported.

4. Click **Import**.

   A report is generated and displayed.

   You can click on the user's name to open the user's personal details and make any required changes. Click the back button in the browser to return to the report.

   New users receive a welcome email that includes a link for creating a password for ArcSight Investigate. The link is valid for 24 hours.

5. If the users or groups in ESM are modified, then you can click **Start New Import** to import again.

   When importing for a second time, changes made in ESM such as deleting a group or removing a user from a group, does not affect the groups and users in ArcSight Investigate.

Managing user groups

This section includes the following topics:

- What is a user group? 13
- Creating a user group 14
- Adding users to a group 14
- Removing a user from a group 15
- Adding or removing a group manager 16
- Deleting a group 16

What is a user group?

A user group is a collection of users managed by the same person. A group can have more than one manager.

Managing users in a group instead of individually, is a practical way for system administrators to delegate managerial capabilities to other administrators in the organization. For example, a system administrator can create several user groups and assign these groups to specific managers. The new managers can manage the users in their groups without having to actually create them.

In addition, managing users in groups provides managerial permissions to specific users only. This is another way of defining a clear scope of capabilities to users and limiting the span of control for admins, other than roles.

Note: Groups and roles are not directly connected. Roles are assigned to a user and not to a group.

New users are automatically added to the All Users group. Users cannot be removed from this group. By default, the system administrator (or system administrators if your organization has more than one) is the manager of this group.

In order to manage a user, that user must belong to one of your groups. Even users that you create yourself must belong to one of your groups in order for you to manage them.

As a group manager, you can perform the following tasks:

- "Managing a user's account" on page 17
- "Adding or removing a group manager" on page 16
- "Terminating processes" on page 24

Tip: If you do not want to create groups, you can make all the admins managers of the All Users group. This way, they can perform all possible tasks on all the users, making groups redundant. For more information, see "Adding or removing a group manager" on page 16.

Creating a user group

Required permissions:

- Manage Groups

When you create a group you automatically become its manager.

Procedure

Location: Left navigation > Admin > User Groups

1. In the top navigation pane, click **Create Group**.
2. In the text box in the title bar, enter a group name, and then press **ENTER**.

   The group is created.
3. To edit a group name, click on the group name to make it editable, and enter a new name.
4. "Adding users to a group" below.

Adding users to a group

Required permissions:

- Assign Users to Groups or Manage Groups

You can add users that you manage to multiple groups. They can be groups that you manage or groups that someone else manages. Once you add a user to a group that someone else manages, that person becomes the user's manager as well. You can also create new users and add them to a group.

You can add users to groups in the following ways:

- From within the group, to that specific group.

  The group that is open in the UI.
- Add users to any group.

  To any group in the system, including ones that you don't manage.
- From the user details page.

Procedure

From within the group:

Location: Left navigation > Admin > User Groups

1. Select the group to which you want to add users.
2. Click **Add Users to this Group**.
3. In the **Add Users to this Group** dialog box, select the users that you want to add to the group, and then click **Add**.

   Only users that you manage (that belong to a group that you manage) are displayed.
4. To add a new user, click **Create New User** and follow the instructions in "Creating a user" on page 10.

From any group:

Location: Left navigation > Admin > User Groups

1. Select the group that includes the users that you want to add to another group. The All Users group includes all the users in the system.
2. Select the users from the list, and then click **Add Users to Another Group**.
3. In the **Add Users to Another Group**, select the group or groups, and then click **OK**.

From the user details page:

1. Search for the user, as described in "Searching for a user" on the next page.
2. Click the **Groups** tab, and then click **Add/Remove User from Groups**.
3. Select the groups to which to add the user, and then click **Save**.

Removing a user from a group

Required permissions:

- Assign Users to Groups or Manage Groups

You can remove users that you manage from any group to which they belong.

If there are users that belong only to this group (aside from the All Users group), we recommend that you move them to another group. Otherwise, these users will not be managed under any group aside from the All Users group.

You can add users to groups in the following ways:

- From the group.
- From the User Details page.

Procedure

From the group:

Location: Left navigation > Admin > User Groups

1. Select the group from which you want to remove the users.
2. Select the users that you want to remove.
3. Click **Remove Users from this Group**.

From the User Details page:

1. Search for the user, as described in "Searching for a user" below.
2. Click the **Groups** tab, and then click **Add/Remove User from Groups**.
3. Select the groups from which you want to remove the user, and then click **Save**.

Adding or removing a group manager

Required permissions:

- Manage Groups

If you are the group manager, you can assign additional managers to a group. Multiple manager for a group can help delegate administrative capabilities, such as managing a user's account.

Procedure

Location: Left navigation > Admin > User Groups

1. Click on the group to which you want to add managers.
2. In the top navigation, click **Add/Remove Group Managers**.
3. In the **Add/Remove Group Managers** dialog box, select or remove managers, and then click **OK**.

Deleting a group

Required permissions:

- Manage Group

You can delete a group if you are a system administrator or if you are a manager of the group.

If there are users that belong only to this group (aside from the All Users group), we recommend that you move them to another group before you delete the group. Otherwise, these users will not be managed under any group aside from the All Users group.

Procedure

Location: Left navigation > Admin > User Groups

1. Select the group that you want to delete.
2. In the top navigation pane, click **Delete Group**.

Searching for a user

Required permission:

- View Users

Procedure

Location: Left navigation > Admin > User Groups

1. In the top navigation, click **Search Users**.
2. In the **Search Users** dialog box, enter one of the following identifiers:
   - Name
   - Email
   - ID
3. Click the user's name to open the **User Details** page.

Managing a user's account

Required permissions:

The required permission depends on the action you want to perform. For example, if you want to reset a user's password, then you must have the Change User Password permission.

You must be the user's manager in order to manage his or her account.

What can you do?

- Edit the user's first name, last name, and email
- Reset the user's password

  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.
- Activate or deactivate a user

  A deactivated user cannot log into the system.
- Unlock a user

  Users are locked after three attempts to log into the system with the wrong credentials. When a user's account is locked, a notification is displayed on the User Details page.

Procedure

Location: Left navigation > Admin > User groups

1. "Searching for a user" on the previous page.
2. To change the user's first name, last name, or email, edit the relevant fields, and then click **Save**.
3. To reset the user's password:
   a. Click **Change Password**.
   b. On the **Change Password** dialog box, enter a new password and confirm it.

c. Click **Save**.

An email is sent to the user with a notification.

4. To activate or deactivate a user, click the **User Status** toggle.

   A deactivated user will not be able to log into the system.

5. To unlock a user account, click **Unlock** in the notification pane. This button is only displayed for locked users.

Managing roles

Managing roles includes creating, editing, and deleting roles. For more information, see "Creating a role" below, "Editing a role" on the next page, and "Deleting a role" on page 20 respectively. To manage roles, you must have the Manage Roles permission.

A user must have at least one role with one functional permission in order to log into the system.

ArcSight Investigate includes built-in roles. To learn more, see "Built-in roles" on page 22.

You can see all your roles and permissions in the following location: My Profile > Roles & Permissions.

What is a role?

A role is a collection of permissions to perform certain operations or to access certain data (fields) in the system. Examples of permissions for performing operations include: View Users, Execute Search, and Manage Roles. Data access means that when you search for events, some fields are visible to you while other may not be. Restricted fields also affect visualizations; you will not be able to view charts that include a restricted field. By default, all fields are accessible. You can restrict fields when you create a role, through the Data Access tab.

Roles are created for various job functions and are assigned to users according to their function in the organization. For example, a level 1 analyst can perform operations on dashboards and searches, but cannot create new users in the system. ArcSight Investigate includes a number of built-in roles. For more information, see "Built-in roles" on page 22.

Permissions cannot be assigned directly to users.

For the full permissions' list, see "Permissions" on page 22.

Creating a role

Required permissions:

- Manage Roles

When you create a new role, it is assigned to you automatically.

To learn more about roles, see "What is a role?" above.

Procedure

Location: Left navigation > Admin > Roles

1. Click **Create Role** in the top right.
2. In the text box in the title bar, enter a name for the role, and then press **ENTER**.
3. To edit a role name, click on the role name to make it editable, and enter a new name.
4. In the **Permissions** tab, select the permissions that you want to assign to this role.

   You can only assign permissions that you have yourself.
5. To restrict data access for this role, do the following:
   a. Click the **Data Access** tab.

      By default, all fields are allowed.
   b. Clear the check box from fields that you want to restrict.

      You can filter fields in the following ways:
      - Click **Restricted Fields Only**. Only restricted fields are displayed.
      - Enter a string or phrase in the **Filter** box.

Editing a role

Required permissions:

- Manage Roles permission, as well as the role itself.

When you edit a role you can add or remove permissions, and add or remove data access. To learn more, see "What is a role?" on the previous page

Note: You cannot edit the System Admin role. For more information, see "Built-in roles" on page 22.

Procedure

Location: Left navigation > Admin > Roles

1. Click the role that you want to edit.
2. To edit the role name, click the role name to make it editable and enter the new name.
3. To edit permissions, in the **Permissions** tab, select or clear the check box of the permissions that you want to add or remove from this role.

   You can only select a permission that you have yourself.
4. To edit data access, click the **Data Access** tab, and then select or clear the check box of the fields that you want to add or remove from this role.

   You can only select fields that you have yourself.

Deleting a role

Required permissions:

- Manage Roles permission, as well as the role itself.

Note: You cannot delete the System Admin role. For more information, see "Built-in roles" on page 22.

Procedure

Location: Left navigation > Admin > Roles

1. Click the role that you want to delete.
2. Click **Delete Role** in the top right.

Assigning a role to a user

Required permissions:

- Any user with the Manage Roles or Assign Roles to Users permission can assign a role to a user in the system. You do not have to manage the user.

You can only assign roles that you have yourself.

Note: If you are creating an Admin user, in addition to the Admin role, make sure to assign the user roles that he can assign to the users that he creates. For example, the Guest and the User roles.

You can assign roles in the following ways:

- From the role page.

  Recommended when you want to assign a specific role to multiple users.
- From the User Details page.

  Recommended when you want to assign multiple roles to a specific user.
- From a specific group.

  Recommended when you want to assign multiple roles to multiple users in a certain group.

Procedure

From the Roles page:

Location: Left navigation > Admin > Roles

1. Click the role that you want to assign.
2. In the role page, click the **Users** tab.
3. In the **Users** tab, click **Assign Role to Users**.

4. In the **Assign Role to Users** dialog box, select the users to which you want to assign the role, and then click **Save**.

From the User Details page:

Location: User Details > Roles & Permissions

1. Search for the user to which you want to add roles, as described in "Searching for a user" on page 16.
2. In the **User Details** page, click the **Roles & Permissions** tab.
3. Click **Assign/Remove Roles**.
4. In the **Assign/Remove Roles** dialog box, select (or clear) the roles for this user, and then click **OK**.

From a specific group:

Location: Left navigation > User Groups > specific group

1. On the group page, select the users to which you want to assign roles.
2. Click **Assign Roles to Users**.
3. Select the roles that you want to add, and then click **Save**.

Removing a role from a user

Required permissions:

- Manage Roles / Assign Roles to Users

You do not have to be the user's manager to remove his or her role; you only need the appropriate permissions. If you are not a group manager, you can only assign or remove roles that you have. If you are a group manager, you can remove any role that the user has but can assign only roles that you have yourself.

You can remove roles in the following ways:

- For a single user
- For multiple users

  Recommended when you want to remove roles from users from the same group.

Procedure

For a single user:

Location: Left navigation > User Management > User Groups

1. "Searching for a user" on page 16.
2. In the **User Details** page, click the **Roles & Permissions** tab.
3. In the **Roles & Permissions** tab, click **Assign/Remove Roles**.
4. Remove roles for this user, and then click **OK**.

Note: You can also assign or remove roles from multiple users from the user group page.

For multiple users:

Location: Left navigation > User Groups > specific group

1. On the group page, select the users from which you want to remove roles.
2. Click **Remove Roles from Users**.

   Note: Only roles that you have yourself are displayed.

3. Select the roles that you want to remove, and then click **Save**.

Built-in roles

ArcSight Investigate includes the following built-in roles:

- System Admin
- Admin
- Analyst L1
- User
- Guest

The System Admin role has the most permissions and the Guest role has the least permissions.

The System Admin role cannot be edited or deleted. If you have the System Admin role you can perform any action in the system, on any user, role, or group. A System Admin is the manager of the All Users group and cannot be removed. A system admin cannot remove his own System Admin role.

Users that have the Manage Roles permission and that have the specific role can edit or delete all the other built-in roles. By default, System Admin and Admin can edit all built-in roles.

Permissions

Note: A user must have at least one role with one functional permission in order to log into the system.

The following table includes the list of permissions.

Permission	Notes
View Users	A user must have this permission in order to access the Admin module. It is included in other permissions, so you do not need to actively select it if you select one of the permissions below.
Create Users	Includes: • View Users • Assign Roles to Users • Assign Users to Groups
Unlock Users	Includes View Users
Activate/Deactivate Users	Includes View Users
Change User Password	Includes View Users
Change User Email	Includes View Users
Assign Roles to Users	Includes View Users
Assign Users to Groups	Includes View Users
Manage Roles	Includes: • View Users • Assign Roles to Users
Manage Groups	Includes: • View Users • Assign Users to Groups
Execute Search	
Export Search Results	Includes Execute Search

Chapter 3: Processes

Processes are searches that users execute. Some searches are complex or are performed on a large amount of data and may take a long time to execute. Multiple searches that are in progress may slow down the performance of ArcSight Investigate. In such cases, an administrator can terminate searches that are in progress and improve performance. For more information, see "Terminating processes" below.

Terminating processes

Required permissions:

- Manage processes

Only processes of users that you manage are displayed.

Procedure

Location

Left navigation > Admin > Processes page

1. Select the desired process and then click **Terminate**.

Chapter 4: Managing dashboard widgets

The Dashboard provides widgets to simultaneously monitor numerous event flows and create text boxes.

Adding a widget to the Dashboard

About

ArcSight Investigate provides for two types of widgets: Data visualization chart and text box. The former is added from the Search page and the latter is created in Dashboard.

Procedure

Location: Left navigation > Dashboard button > Dashboard page

1. To add a text box widget to the Dashboard, click **Add Text Box** in the task bar.
   - Type text directly in the text box.
   - To delete the text box, choose ••• > **Delete** from the text box panel.

Location: Left navigation > Search button > Search page

2. To add a chart to the Dashboard, open the desired search and then expand the "Visualize" area.
3. Choose ••• > **Add to Dashboard** from the chart panel of the desired chart.
   - Repeat this step for each chart that you want to add to the Dashboard.
   - To change the chart before adding it to the Dashboard, see "Editing search-results charts" on page 48
   - To delete the chart from the Dashboard, in the Dashboard choose ••• > **Delete** from the chart panel.

Location: Left navigation > Dashboard button > Dashboard page

4. To view the latest data for a chart associated with a real-time search, choose ••• > **Refresh** from the chart panel.

   Refresh a chart associated with a real-time search in order to see the latest data in the chart.

See Also

- "Creating line, bar, column, and area charts" on page 39
- "Creating a pie chart" on page 44
- "Creating a scatter plot chart" on page 46

Deleting a widget from the Dashboard

About

ArcSight Investigate provides for two types of widgets: Data visualization chart and text box. The former is added from the Search page and the latter is created in Dashboard.

Procedure

Location: Left navigation > Dashboard button > Dashboard page

- For the desired widget, choose ••• > **Delete**.

  A deleted data visualization chart is removed from the Dashboard, but not from the originating search in the Search page.

See Also

- "Editing search-results charts" on page 48

Chapter 5: Searching event data

ArcSight Investigate helps you search event data as part of an ad hoc analysis and investigation of security incidents and threat context. The context of a search is twofold:

- Investigate an alert or incident from ArcSight Enterprise Security Manager (ESM) or ArcSight Command Center (ACC). Or, an investigative search can originate from a manual notification regarding a particular entity such as a user, server, or IP address.
- Gather and compile data for analysis and reporting on large-scale (billions of events) data. From a generated sparkline, chart, and table of events, you can detect anomalies in search responses that point to a security threat.

Searching events

About

- The query can either be full text, natural language, or contextual.
- An event search consists of specifying query input, search-result fields, and the time period to search within.
- You can filter search results for more specific information.
- ArcSight Investigate supports up to 10 session searches and up to 40 saved searches.
- If an event does not have data for a schema field, it is represented as null [***(null)***, *(nuLL)*] in the Search page, except in the Events table where an empty cell is used.

Procedure

Location: Left navigation > Search button > Search page

1. Click **New Search** in the left navigation area.
   - The search appears in a new tab in the left navigation.
   - To use an existing search:
     - Session search — From the left navigation, open the desired search from the list of session searches.

- Saved search — From the left navigation, click **Saved Results** and open the desired search.

To modify any of these searches, continue with this procedure.

2. Accept the default search name or rename the search.

   Investigate automatically assigns a name to a new search in the format of `Search x + 1`.

3. To search for fields with data, specify the desired query input from the **Search** field.

   - A drop-down appears below the field suggesting applicable information.
   - You can enter up to 20 search items, which are not case sensitive.
   - You can copy and paste in the **Search** field.
   - Investigate auto completes search items based on a schema data dictionary.
   - To use "canned" queries, type # and then select the desired query.
     - There are eight "canned" queries, indicated by the **preset** label.
     - To view the query input for the "canned" query, hover over the "canned" query in the **Search** field.
   - For multiple search items, Investigate inserts an `or` operator between each item.
   - You can enter multiple values for a search item.

Examples

Multiple values for search item

```
IP = 15.214.133.85, 15.211.201.91, 15.215.101.77, 15.218.151.87,
15.212.145.22
```

Multiple types of search items

```
IP = 15.214.133.85, 15.211.201.91, 15.215.101.77, 15.218.151.87 and agt
= 15.214.130.65 and ahost = n15-214-130-h65.arst.usa.hp.com
```

- You can specify fields that are not in the standard categories.
- The query input determines the search type (full text, natural language, or contextual).
- To remove a search item, use the backspace key.

4. To search for a field without data, use the `null` field value in the **Search** field.

   - The `Null` and `NULL` query formats are also supported.
   - If an event does not have data for a schema field, it is represented as null [***(null)***, *(`nuLL`)*] in the Search page, except in the Events table where an empty cell is used.

5. Accept the default fieldset for search result events, or change the fieldset by clicking **Default Fieldset**.

   - At anytime during the search procedure, you can change the fieldset.

     Any existing data visualization charts adjust to the fieldset change, along with the Events table.
   - Depending on your data access permissions, you may not see all the possible fields for an event.

     From the user page (**Left navigation > [user] > My Profile > Data Access > My Profile**), you can view event fields for which you have access.

   See

   - "Managing search-results fieldsets" on page 35
   - "Creating a fieldset for search results" on page 35
   - "Editing a fieldset for search results " on page 36
   - "Deleting a fieldset for search results " on page 37

6. From the time drop-down accept the default time range (**Last 30 minutes**), or specify a different time range to be searched.

   - The **Quick Range** feature provides a convenient way to select common search periods ranging from the past minute up to a year ago.

- A **Quick Range** selection appears in the **Custom Range** fields, where you can edit this information.
- By default, the **Custom Range** fields contain the **Quick Range** time range for the last 30 minutes (current time minus 30 minutes).

7. Click **Search**.
   - If results are found, "Timeline" and the Events table are populated with data. Also, the **Filter** button and **Create Visualization** button appear.
   - To cancel the search, click **X** in the **Search** field.

     This will not remove the search conditions that you originally specified.
   - To change the search, edit the query input in the **Search** field and then click **Search**.

     The original search stops and the new query begins.
8. To filter the search results:
   a. Click **Filter** and then make the appropriate selections from the **Select** drop-downs.

      To add another filter, click **Add Filter**.

      Other filtering methods:
      - Drag and drop a statement from the Events table to the "Filter" area.

        Upon drop, the "Filter" area highlights if the statement is acceptable.
      - Right click a data cell in the Events table and then choose **Use As A Filter**.

      For both cases, Investigate adds a new filter (row) to the "Filter" area.

   b. Click **Apply** or **Apply all**, depending on your filter usage.
      - The updated search results are reflected in "Timeline", any data visualizations, and Events table.
      - Investigate list filters in abbreviated form.

        Examples

        ```
        IP = "208.202.19.230"
        Username = bob
        Rt = 1445077926551
        ```

        To see the whole filter, click the desired abbreviation.
9. To view search results graphically, expand the "Visualization" area and then create the desired chart (see "Charting search-results data" on page 38).
10. To manage search results information, expand the "Events" area and specify how event data should appear (see "Managing search results information" on page 51).

11. Click **Save** from the task bar.
    - The query input and results are saved along with any data visualization charts you created.
    - Investigate supports up to 10 session searches and up to 40 saved searches.
    - If you reach the limit of session searches, you will have to delete one in order to create a new search.
    - Depending on your data access permissions, you may not see all the possible fields for an event.

      From the user page (**Left navigation > [user] > My Profile > Data Access > My Profile**), you can view event fields for which you have access.
12. To do another search, either create a new one or open an existing one.
    - An existing search can be either a session or saved search
    - Investigate supports up to 10 session searches and up to 40 saved searches.

      If total session searches is at the limit of 10, delete a session searches in order to create a new search.
    - You can open or create a new search before the current one finishes.
    - To complete a new search, repeat steps 1 — 11.

See Also

- "Searching events in ESM" below
- "Searching on null event fields" on page 1
- "Viewing search results for a time range" on page 38
- "Charting search-results data" on page 38
- Creating search-results charts
- "Editing search-results charts" on page 48
- "Deleting search-result charts" on page 50

Searching events in ESM

About

- ArcSight Enterprise Security Manager (ESM) (Console or ArcSight Command Center (ACC)) enables you to investigate events from a channel being analyzed in ArcSight Investigate.
- Within ESM, you can searches on a maximum of five fields.

  Within Investigate, you can filter ESM data for more specific information.
- See the ESM 6.11 guide for Console and ACC and ESM 6.11 release notes guide

Procedure

Prerequisite

- New ESM user:

  Run the configuration wizard (see "Using the Configuration Wizard" in the *HPE Security ArcSight ESM Installation Guide*).

- Existing ESM user who upgraded to 6.11:

  a. Login to the system and go to `/etc/bin` directory
  b. Stop the manager.

     ```
     ./arcsight_services stop manager
     ```

  c. `cd /opt/arcsight/manager/bin`
  d. `./arcsight managersetup`
  e. Enable Investigate as done in the configuration wizard for a fresh install.
  f. Start the manager.

     ```
     ./arcsight start manager
     ```

ESM Console Location:	ESM > Event Detail (Inspect/Edit) panel
	ESM > Active Channel
ACC Location:	ESM > Event Detail
	ESM > Active Channel
	ESM > Visualize Events
	ESM > Active Channel
	ESM > Dashboards

1. Open an event viewer such as an active channel, or view event details in the Inspect/Edit panel.
2. Right-click a row and make the appropriate selection.

   Note: ESM fields that are not supported in Investigate searches appear disabled.

   For a list of supported fields, see "Running ArcSight Investigate Searches" under "Usage Notes" in the ESM 6.11 Release Notes. Also see the 6.11 *HPE Security ArcSight ESM Command Center User's Guide*.

   - To search on a specific value, select **ArcSight Investigate**.

     ESM generates a URL comprised of:

     - `query` equals `SELECTED COLUMN NAME = <title of column>`

       Example

       ```
       SELECTED COLUMN NAME = Employee
       ```

- `StartTime` equals `<time stamp>`
- `endTime` equals `<time stamp>`

- To search on one or more values, select **ArcSight Investigate (Multiple Fields)**.
  - From the dialog, select up to five fields for investigation. These fields are based on the columns that are available on the channel.
  - ESM generates a URL comprised with:
    - `query` equals `SELECTED COLUMN NAME = <title of column>`

      Example

      `SELECTED COLUMN NAME = Employee`
    - `StartTime` equals `<time stamp>`
    - `endTime` equals `<time stamp>`
- ESM opens a browser with the ESM-generated URL and creates a new search in Investigate.

  To modify this search, continue with this procedure.

3. Accept the default search name or rename the search.

   Investigate automatically assigns a name to a new search in the format of `Search x + 1`.

4. Accept the default fieldset for search result events, or change the fieldset by clicking **Default Fieldset**.

   - At anytime during the search procedure, you can change the fieldset.

     Any existing data visualization charts adjust to the fieldset change, along with the Events table.
   - Depending on your data access permissions, you may not see all the possible fields for an event.

     From the user page (**Left navigation > [user] > My Profile > Data Access > My Profile**), you can view event fields for which you have access.

   See

   - "Managing search-results fieldsets" on page 35
   - "Creating a fieldset for search results" on page 35
   - "Editing a fieldset for search results " on page 36
   - "Deleting a fieldset for search results " on page 37

5. From the time drop-down, accept the time range or specify a different one to be searched.

   - An investigation done with integration commands results in the default time range (**Last 30 minutes**) being used.

- An investigation originating from an active channel results in the channel time range being used.
- The Quick Range feature provides a convenient way to select common search periods ranging from the past minute up to a year ago.
- A Quick Range selection appears in the **Custom Range** fields, where you can edit this information.
- By default, the **Custom Range** fields contain the **Quick Range** time range for the last 30 minutes (current time minus 30 minutes).

6. Click **Search**.

   - If results are found, "Timeline", Events table, and any data visualization are populated with data. Also, the **Filter** button appears.
   - If a field from ESM is empty, such as `name = ''`, it appears in the Investigate **Search** field as `name = '', null`. The query searches for both `''` and `null` values.

     Note: If a numeric field from ESM is empty, then remove `''` before executing the query in order to prevent an error.

     Example: If ESM launches Investigate with `Bytes In = ''`, the query translates to `Bytes In = '', null` in the Investigate **Search** field. Change the query to `Bytes In = null` before executing the query.

   - To cancel the search, click **X** in the **Search** field.
     This will not remove the search conditions that you originally specified.
   - To change the search, edit the query input in the **Search** field and then click **Search**.
     - The original search stops and the new query begins.
     - To use "canned" queries, type # and then select the desired query.
       - There are eight "canned" queries, indicated by the **preset** label.
       - To view the query input for the "canned" query, hover over the "canned" query in the **Search** field.

7. To filter the search results:
   a. Click **Filter** and then make the appropriate selections from the **Select** drop-downs.

      To add another filter, click **Add Filter**.
      Other filtering methods:
      - Drag and drop a statement from the Events table to the "Filter" area.
        Upon drop, the "Filter" area highlights if the statement is acceptable.
      - Right click a data cell in the Events table and then choose **Use As A Filter**.

For both cases, Investigate adds a new filter (row) to the "Filter" area.

b. Click **Apply** or **Apply all**, depending on your filter usage.

- The updated search results are reflected in "Timeline", any data visualizations, and Events table.
- Investigate list filters in abbreviated form.

  Examples

  ```
  IP = "208.202.29.230"
  Username = bob
  Rt = 1445077926551
  ```

  To see the whole filter, click the desired abbreviation.

8. Click **Save** from the task bar.

- The query input and results are saved along with any filter, data visualization, and time range selection.
- Depending on your data access permissions, you may not see all the possible fields for an event. From the user page (**Left navigation > My Profile > Data Access > My Profile**), you can view event fields for which you have access.

See Also

- "Searching events" on page 27
- "Viewing search results for a time range" on page 38
- "Charting search-results data" on page 38
- Creating search-results charts
- "Editing search-results charts" on page 48
- "Deleting search-result charts" on page 50

Managing search-results fieldsets

The default fieldset contains the 63 most common event fields. These fields are available for creating data visualization charts and for viewing in the Events table. Each field (column) in this table can provide the 10 most and least common values.

Creating a fieldset for search results

About

The fieldset determines what search results fields are available for creating data visualization charts and

what data columns appear in the Events table.

Procedure

Location: Left navigation > Search button > Search page

1. Click the fieldset button.
   - For a new search, the button defaults to **Default Fieldset**.
   - For an existing search, the button name is that of the last selected fieldset, or the default fieldset if no selection was made.
2. From the Show Fields dialog, select **Create a new set** from the drop down.
3. Select and deselect any necessary fields.

   To quickly locate a desired field, use the **Filter** field.
4. Accept the default fieldset name or rename it.
5. Click **Save**.

See Also

- "Editing a fieldset for search results " below
- "Deleting a fieldset for search results " on the next page

Editing a fieldset for search results

About

The fieldset determines what search results fields are available for creating data visualization charts and what data columns appear in the Events table.

There are a possible of 176 fields in a fieldset. To manage the display of this data, you may want to limit the number of fields in a fieldset.

Procedure

Location: Left navigation > Search button > Search page

1. Click the fieldset button.

   The button name is that of the last selected fieldset, or the default fieldset if no selection was made.
2. To select a different fieldset, use the drop-down in the Show Fields dialog.
3. From the drop-down, select **Edit this set**.
4. Specify the desired fields.

- To quickly locate a field, use the **Filter** field.
- To conveniently manage the selection of fields in a fieldset, use the **Select all**, **Unselect all**, and **View all** buttons.

> **Note:** If you remove a field from a fieldset, Investigate will remove all the filters, queries, and charts that use that field.

5. To change the fieldset name, use the **Fieldset Name** field.
6. Click **Save**.

See Also

- "Creating a fieldset for search results" on page 35
- "Deleting a fieldset for search results " below

Deleting a fieldset for search results

About

- The fieldset determines what search results fields are available for creating data visualization charts and what data columns appear in the Events table.
- If only a single search is using a particular fieldset, then the fieldset can be deleted and the default fieldset will be used in its place.
- If two or more searches are using a particular fieldset, then the fieldset cannot be deleted.

Procedure

Location: Left navigation > Search button > Search page

1. Click the fieldset button.

   The button name is that of the last selected fieldset, or the default fieldset if no selection was made.

2. To select a different fieldset, use the drop-down in the Show Fields dialog.
3. From the drop down, select **Edit this set**.
4. Click **Delete**, and then **Save**.

See Also

- "Creating a fieldset for search results" on page 35
- "Editing a fieldset for search results " on the previous page

Viewing search results for a time range

About

- A Timeline selection focuses on the search results for a specific time range within the specified time period.
- The results of a Timeline selection are reflected in any data visualization and the Events table.

Procedure

Prerequisite

"Searching events" on page 27

Location: Left navigation > Search button > Search page > Timeline area

1. Ensure that the Timeline chart is visible.
2. Ensure that the **Range Selector** is enabled.

   With the range selector off, you can hover over a sparkline and view the specific time range and number of events for the sparkline.

3. Position the pointer in the time line and using the double-arrow cursor, drag in the desired time range.

   - The time scale for a large specified time range may have large increments.
   - To aid in your selection, the exact desired time range appears above the chart.
   - The event counts for the specified time period and selected time range appear above the chart.
   - If you turn off the **Range Selector**, and turn it back on again, the specified time range will still be selected.

4. To remove a specified time range, click **X** at the end of the "Timeline" time range.
5. To save the time range selection, click **Save**.

   The query input is also saved along with any filter and data visualization.

See Also

- "Visualizing search-results data" on page 1

Charting search-results data

To better understand search-results data, you can represent it graphically. From the Search page, data visualization enables you to add up to 10 charts.

ArcSight Investigate provides data comparison charts and non-comparisons data charts. Data comparison charts include line, column, bar, and area. Non-comparisons data charts include pie and scatter plot.

Charts (widgets) can be added from the Search page to the Dashboard (see "Adding a widget to the Dashboard" on page 25).

Creating line, bar, column, and area charts

About

- Line, bar, column, area and scatter plot are data-comparison charts. For these charts, you can create up to six series of data comparisons.
- The first chart series sets the X- and Y-axis parameters, which remain set for any subsequent series.
- Any ordering specified in the first chart series applies to any subsequent series.
- For any subsequent chart series, you can specify different fields for **Filter By** and **Order By** and set aggregate functions for these parameters along with X- and Y-axis parameters.
- Parameters that can accept the field highlight.
- Parameters that cannot accept the field turn red.
- X- and Y-axis options

Field type	X-axis function	Y-axis function
Time	minute hour day (default) week month year time value itself	count Example: Count the number of events for the time period.

String	String value itself (default) len	count (default) count distinct len
Number	actual number (default)	count count distinct sum (default) average max min Number value itself (only for scatter plot) **Note:** For the average function, the default is regular average. Example: For bytes out, the average will be sum (Bytes Out) / number of events (that contains bytes out). If you selected **Group By User**, then ArcSight Investigate uses the formula: (sum (Bytes Out (only for events when `user!=Null`)) / distinct number of users (without `Null`)

- When dragging a continuous-value field to a discrete-value parameter, Investigate applies an aggregate function to the field.

  Example

  For **Event Time**, Investigate applies the aggregate function, `day(Event Time)`.

- When dragging a discrete-value field to a continuous-value parameter, Investigate converts the field to a continuous-value field.

  Example

  For **File Name**, Investigate applies the applies the `count()` function.

- Within a parameter, fields appear in the following formats:
  - Single key/value pair — `<field>:<value>`

    Example

    ```
    department:sales
    ```

  - Single key with multiple values — `<field>:<value1>,<value2>,...`

    Example

    ```
    user:johnny, bob,...
    ```

- Aggregate function — `<function>(<field>)`

  Example

  ```
  - sum(Bytes Out)
  - month(Event Time)
  ```

Procedure

Prerequisite

"Searching events" on page 27

Location: Left navigation > Search button > Search page > Visualize area

1. Expand the "Visualize" area.
2. Click **Create Visualizations**.
3. From the Choose Visualizations dialog, select the desired data comparison chart type.

   - Line, bar, column, and area are data-comparison charts.

4. To change available fields for parameter selection, click **Default Fieldset** (see "Editing a fieldset for search results " on page 36).
5. Drag the desired field to **X-Axis**.

   - The parameter can receive a field with a continuous value .
   - Investigate applies the `sum()` aggregate function to any continuous-value field.
   - Investigate convert a discrete-value field to a continuous value by applying the `count()` aggregate function to the field.
   - The field specified for this parameter remains for any subsequent chart series. You can change the aggregate function for the parameter in a subsequent chart series.

   To change the aggregate function for the parameter

   a. Click the field in **X-Axis**.
   b. From the X-Axis dialog, select the desired value type from the **Aggregate values using** drop down.

      Only aggregate functions that can be applied to continuous-value fields are enabled.

6. Drag the desired field to **Y-Axis**.

   - The parameter can receive any discrete-value field.
   - Investigate applies the `count()` aggregate function to any continuous-value fields.
   - When dragging a field to **Y-Axis**, the scale and data labels of the axis appear.

The label for **Y-Axis** is in the form, `<function>(<field>) or <field>`.

- The field specified for this parameter remains for any subsequent chart series. You can change the aggregate function for the parameter in a subsequent chart series.

To change the aggregate function for the parameter

a. Click the field in **Y-Axis**.
b. From the Y-Axis dialog, select the desired value type from the **Aggregate values using** drop-down.

   Only aggregate functions that can be applied to discrete-value fields are enabled.

7. To compare a data segment against other search results data, drag the desired field to **Filter By**.

   - The parameter can receive multiple discrete fields.
   - The parameter is enabled when both the X- and Y-Axis parameter are set.
   - By default, only some of the values for the field are used.

   To change field values for filtering

   a. Click the field in **Filter By**.
   b. From the Filter By dialog, specify desired field values for filtering.

      Select all and **Unselect all** are convenient for a field with many values where you want to selectively pick values.

8. To specify the field by which records should be ordered, click **Order by**.

   - The parameter is enabled when the **X-Axis** parameter is set.
   - The parameter orders the bars in the Y-Axis
   - Sort orders are dependent on the field used for the Y-Axis.
   - Records appear in ascending order by default.

9. For a Stacked Bar Chart, specify any segmenting of Y-Axis bars by dragging the desired field to **sub-categories**.

   A sub-category enables you to specify a secondary discrete-value field. Each bar in the Y-axis is segmented by this secondary category.

10. To view values in the chart, choose **Preview > Data Labels**.
11. To set a baseline for which to compare chart data, choose **Preview > Plot Line** and then specify baseline value in the adjacent field.

A dotted red line appears in the chart at the specified value .

12. To create another data segment comparison, click **Add Series** and specify any new parameters and aggregate functions.

    - The **X-Axis** and **Y-Axis** parameters specified in the first chart series are inherited in this series.
    - The aggregate functions for parameters in the first chart series are inherited in this series.
    - To hide a chart object (line, bar, or column), click the color of the chart object in the chart legend. Undo this hide by clicking on the appropriate grayed color in the chart legend. This feature is especially useful if a chart object overlaps another.

    To specify data for a subsequent series chart

    a. To change the aggregate function for **X-Axis**, click in this parameter and then make the appropriate changes in the X-Axis dialog (see "To change the aggregate function for the parameter" on page 41).

    b. To change the aggregate function for **Y-Axis**, click in this parameter and then make the appropriate changes in the Y-Axis dialog (see "To change the aggregate function for the parameter" on the previous page).

    c. Drag the desired field to **Filter By**.

    If you change **Order By** for any series, all chart series are affected.

    d. For a bar chart, to segment Y-Axis bars by a secondary category, see step 10.

13. To specify the number of data points for a chart, choose ... > **Points to plot** from the chart panel.

    - The default number of data points is 15, potentially allowing for multiple chart pages for a data-intense chart.
    - If the chart is data intense, you can increase the number of data points in order to reduce the chart page count. The more data points that you specify, the more condense the chart, resulting in less pages.
    - If you add the chart to the Dashboard (see next step), the chart in the Dashboard will not reflect the data-point count that you specify here.

14. To add the chart to the Dashboard, choose ... > **Add to Dashboard** from the chart panel.

    - It is not necessary to save the search before adding the chart to the Dashboard
    - Investigate places the chart in the last position in the Dashboard.

To remove the chart from the Dashboard

- Choose ••• > **Delete** from the chart panel.

15. To enlarge the chart, click ↗.
    - Any other charts are hidden until your return the chart to its original size.

    To return the chart to its original size

    - Click the left-pointing arrow in the chart title bar.

16. To name (rename) the chart, choose ••• > **Rename** from the chart panel.
    - The default name is the **X-Asis** parameter field name combined with the **Y-axis** parameter field name.

      Example

      BYTES OUT BY AGENT ADDRESS, AGENT SEVERITY BY DESTINATION USERNAME

    - If you previously added the chart to the Dashboard, the new chart name will also appear for that chart.

17. Click **Done** and then **Save** in the task bar.
18. To add another chart, click ✚ and repeat steps 3 - 18.

Creating a pie chart

About

- Charts (widgets), including a pie chart, can be added from the Search page to the Dashboard.
- Parameters that can accept a field highlight.
- Parameters that cannot accept a field turn red.

Procedure

Location: Left navigation > Search button > Search page > Visualize area

1. Expand the "Visualize" area.
2. Click **Create Visualizations**.
3. From the Choose Visualizations dialog, select **PIE CHART**.

   To use a chart saved as a template

   - Click **Choose Template**.

4. Drag the desired field to **Measure**.

- This parameter takes a continuous-value field.
- The field-supported aggregate function is applied.
- Discrete-value fields are automatically converted to a continuous-value field with the `count()` aggregation function.
- The parameter determines the size of the pie slice.

To change the aggregate function for the parameter

a. Click the field in **Measure**.
b. From the Measure dialog, select the desired value type from the **Aggregate values using** drop down.

   Only aggregate functions that can be applied to continuous-value fields are enabled.

5. Drag the desired field to **Label**.

- This parameter takes a discrete-value field. Events are grouped by unique values for the specified field.

To change the aggregate function for the parameter

a. Click the field in **Label**.
b. From the Label dialog, select the desired value type from the **Aggregate values using** drop down.

   Only aggregate functions that can be applied to discrete-value fields are enabled.

6. Drag the desired field to **Filter By**.

   The parameter can receive multiple Discrete-value fields.

To change the aggregate function for the parameter

a. Click the field in **Filter By**.
b. From the Filter By dialog, deselect unwanted field values.

7. To add the chart to the Dashboard, choose ••• > **Add to Dashboard** from the chart panel.

- It is not necessary to save the search before adding the chart to the Dashboard
- Investigate places the chart in the last position in the Dashboard.

To remove the chart from the Dashboard

- Choose ••• > **Delete** from the chart panel.

8. To enlarge the chart, click ↙↗.
   - Any other charts are hidden until your return the chart to its original size.

   To return the chart to its original size

   - Click the left-pointing arrow in the chart title bar.

9. To name (rename) the chart, choose ••• > **Rename** from the chart panel.

   If you previously added the chart to the Dashboard, the new chart name will not appear in the Dashboard.

10. Click **Done** and then **Save** in the task bar.
11. To add another chart, click ✚ and repeat steps 3 - 11.

See Also

- "Creating line, bar, column, and area charts" on page 39
- "Creating a scatter plot chart" below
- "Adding a widget to the Dashboard" on page 25

Creating a scatter plot chart

About

Charts (widgets), including a scatter plot chart can be added from the Search page to the Dashboard.

- Parameters that can accept a field highlight.
- Parameters that cannot accept a field turn red.

Procedure

Location: Left navigation > Search button > Search page > Visualize area

1. Expand the "Visualize" area.
2. Click **Create Visualizations**.
3. From the Choose Visualizations dialog, select **SCATTER PLOT**.
4. Drag the desired field to **X-Axis**.
   - The parameter can receive a field with a continuous value.
   - Investigate applies a field-supported aggregate function to any continuous-value field.
   - Investigate convert a discrete-value field to a continuous value by applying the `count()` aggregate function to the field.

5. Drag the desired field to **Y-Axis**.
   - The parameter can receive any discrete-value field.
   - Investigate applies the `count()` aggregate function to any continuous-value fields.

   To change the aggregate function for the parameter

   a. Click the field in **Y-Axis**.
   b. From the Y-Axis dialog, select the desired value type from the **Aggregate values using** drop down.

      Only aggregate functions that can be applied to discrete-value fields are enabled.

6. When dragging a field to **Y-axis**, the scale and data labels of the axis appear.

   The label for **Y-axis** is in the form, `<function>(<field>) or <field>`.
7. Drag the desired field to **Category**.
   - This parameter takes a discrete-value field.
   - Each unique value for that field is represented by a different color point.

8. To add the chart to the Dashboard, choose ••• > **Add to Dashboard** from the chart panel.
   - It is not necessary to save the search before adding the chart to the Dashboard
   - Investigate places the chart in the last position in the Dashboard.

   To remove the chart from the Dashboard

   - Choose ••• > **Delete** from the chart panel.

9. To enlarge the chart, click ↗↙.
   - Any other charts are hidden until your return the chart to its original size.

   To return the chart to its original size

   - Click the left-pointing arrow in the chart title bar.

10. To name (rename) the chart, choose ••• > **Rename** from the chart panel.

    If you previously added the chart to the Dashboard, the new chart name will not appear in the Dashboard.

11. Click **Done** and then **Save** in the task bar.
12. To add another chart, click ✚ and repeat steps 3 - 12.

See Also

- "Creating line, bar, column, and area charts" on page 39
- "Creating a pie chart" on page 44
- "Adding a widget to the Dashboard" on page 25

Editing search-results charts

About

Charts (widgets) can be added from the Search page to the Dashboard.

Procedure

Location: Left navigation > Search button > Search page > Visualize area

1. Ensure that the "Visualize" area is expanded.
2. To rename the chart, choose ... > **Rename** from the chart panel.
   - The default name is the **X-asis** parameter field name combined with the **Y-axis** parameter field name.

     Example

     BYTES OUT BY AGENT ADDRESS, AGENT SEVERITY BY DESTINATION USERNAME
   - If you previously added the chart to the Dashboard, the new chart name will also appear for that chart.
3. To expand a chart for editing, click and then **Edit**.

   To return to the normal viewing mode, click the left-pointing arrow in the upper left corner.
4. To replace a parameter field:
   a. Hover over the parameter and then click **X**.
   b. Drag the new field to the parameter.
      - Parameters that can accept a field highlight.
      - Parameters that cannot accept a field turn red.
5. To change the aggregate function for the **X-Axis**, **Y-Axis**, **Measure**, and **Label** parameter fields:
   a. Click the parameter field.
   b. From the dialog, select the desired value type from the **Aggregate values using** drop-down.
   c. Click **Close**.
6. To change field values for filtering:

a. Click the field in **Filter By**.
b. From the Filter By dialog, specify desired field values for filtering.

   Select all and **Unselect all** are convenient for a field with many values where you want to selectively pick values.

7. To change the field by which records should be ordered, click the **Order By** parameter field.

   - This parameter field applies only to data-comparison charts (Line, bar, column, and area).
   - The parameter is enabled when both **X** and **Y-Axis** parameters are set.
   - The parameter orders the bars in the Y-Axis
   - Sort orders are dependent on the field used for the Y-Axis.
   - Records appear in ascending order by default.

8. To change the segmenting of Y-Axis bars in a bar chart:

   a. Hover over the **sub-categories** parameter field and click **X**.
   b. Drag the desired field to the parameter field.

   A sub-category enables you to specify a secondary discrete-value field. Each bar in the Y-axis is segmented by this secondary category.

9. To view values in the chart, choose **Preview > Data Labels**.
10. To set a baseline for which to compare chart data, choose **Preview > Plot Line** and then specify baseline value in the adjacent field.

    A broken red line appears in the chart at the specified value .

11. To create another data segment comparison for data-comparison charts, click **Add Series** and specify any new parameters and aggregate functions.

    - The **X-Axis** and **Y-Axis** parameters specified in the first chart series are inherited in this series.
    - The aggregate functions for parameters in the first chart series are inherited in this series.
    - To hide a chart object (line, bar, or column), click the color of the chart object in the chart legend. Undo this hide by clicking on the appropriate grayed color in the chart legend. This feature is especially useful if a chart object overlaps another.

    To specify data for a subsequent series chart

    a. To change the aggregate function for **X-Axis**, click in this parameter and then make the appropriate changes in the X-Axis dialog (see step 5).
    b. To change the aggregate function for **Y-Axis**, click in this parameter and then make the appropriate changes in the Y-Axis dialog (see step 5).

c. Drag the desired field to **Filter By**.

If you change **Order By** for any series, all chart series are affected.

d. For a bar chart, to segment Y-Axis bars by a secondary category, see step 8.

12. To specify the number of data points for a chart, choose ... > **Points to plot** from the chart panel.

- The default number of data points is 15, potentially allowing for multiple chart pages for a data-intense chart.
- If the chart is data intense, you can increase the number of data points in order to reduce the chart page count. The more data points that you specify, the more condense the chart, resulting in less pages.
- If you add the chart to the Dashboard (see next step), the chart in the Dashboard will not reflect the data-point count that you specify here.

13. To add the chart to the Dashboard, choose ... > **Add to Dashboard** from the chart panel.

- Investigate places the chart in the last position in the Dashboard.

To remove the chart from the Dashboard

- Choose ... > **Delete** from the chart panel.

14. Click **Done** and then **Save** in the task bar.

See Also

- "Charting search-results data" on page 38
- "Charting search-results data" on page 38
- "Charting search-results data" on page 38
- "Charting search-results data" on page 38

Deleting search-result charts

Procedure

Location: Left navigation > Search button > Search page > Visualize area

1. Ensure that the "Visualize" area is expanded.

2. For the desired chart, choose ... > **Delete**.

If you added the chart to the Dashboard, the deleted chart will remain in the Dashboard.

- To delete the chart from the Dashboard, see "Deleting a widget from the Dashboard" on page 26.
- If you delete a chart from the Dashboard, the deleted chart is not removed from the search in the Search page.

Managing search results information

ArcSight Investigate lists search-result events in the Events table. The various fields of the event records are represented by the table column headers. Field columns are determined by the fieldset being used and the Show Columns feature.

To manage data in the Events table, Investigate enables you to view the most and least common values for an event record field, pin field columns to better compare values, view all fields of an event, and view select event data.

Viewing the most and least common values for an event record field

About

To help you filter for data security threats, ArcSight Investigate enables you to quickly display the most and least common values for a field. This translates into the count and percentage of hits for the value.

For example, the **devicevendor** field could have a top value of "bluecoat" with a count of 3,000 hits, which is 30% of 10,000 results.

Procedure

Prerequisite

"Searching events" on page 27

Location

Left navigation > Search button > Search page > Events area

1. Ensure that the Events table is visible.
2. From the desired field column header, right click and then select **Preview Top/Least**.

   The default view in the dialog is the greatest data values in descending order.

3. To view the least common data values, click **View bottom 10** from the dialog.

See Also

- "Viewing all fields of an event" on the next page
- "Viewing select event data" on page 53

Pinning field columns to help analyze events

About

It can be helpful to pin a field column—to make a column horizontally stationary, in order to better compare the column values against those of other columns.

Procedure

Prerequisite

"Searching events" on page 27

Location

Left navigation > Search button > Search page > Events area

1. Ensure that the Events table is visible.
2. From the desired column header, right click and select **Make Sticky**.
   - The column moves to the extreme right in the Events table.
   - The column will not scroll horizontally.
   - You can make multiple columns stationary.
   - If there is an existing stationary column, the new stationary column is positioned to the right of that one.
3. To release a stationary column, right click the column header and then select **Unstick**.

See Also

"Hiding a field column to help analyze field values" on page 1

Viewing all fields of an event

About

- ArcSight Investigate enables you to drill into an event record and view all the record fields. Using this feature, you can quickly view the details of a single event without having to add all the fields of the fieldset.
- Depending on your data access permissions, you may not see all the possible fields for an event.

  From the user page (**Left navigation > [user] > My Profile > Data Access > My Profile**), you can view event fields for which you have access.

Procedure

Prerequisite

"Searching events" on page 27

Location

Left navigation > Search button > Search page > Events area

1. Ensure that the Events table is visible.
2. Click the arrow of the desired agent address.

   All the field information for the event record displays, not just that from the specified fieldset.

3. To view information about a field, click the arrow of the desired field.

   You can have multiple fields open at the same time.

See Also

"Viewing the most and least common values for an event record field" on page 51

Viewing select event data

About

From Events table, you can specify how event-record data is displayed.

- View only a select event or events
- Limit the display of field columns
- Sort values for an event field
- Arrange event records based on a common field value

Procedure

Prerequisite

"Searching events" on page 27

Location: Left navigation > Search button > Search page > Events area

1. Ensure that the Events table is visible.
2. To view a selected event, click in the event row and then click **Only show selection**.

   Use the shift and control keys to select multiple events.

3. To limit the display of field columns, click and then deselect undesired fields from the Show Columns dialog.

- When using a fieldset of more than 20 fields, it may be easier to quickly hide and reveal field columns, rather than scrolling left and right, or rearranging the column order.
- To find a field, you can use the filter feature.
- To directly hide a field column in the Events table, right click the desired column header and then select **Hide Column**.
- To reveal a hidden field column in the Events table, click and then select the hidden field column.

4. To reorder field columns, click and drag desired columns to new positions.
5. To sort values for a field column, click the appropriate arrow in the desired column heading.
   - The default order is descending.
   - Single click the column heading for ascending order and double click the column heading for descending order.
6. To view all the event data based on a particular field value, right click on the desired field value and then select **Search For**.

   Below the selected event record, Investigate lists all the event records with a matching field value.

See Also

- "Viewing the most and least common values for an event record field" on page 51
- "Hiding a field column to help analyze field values" on page 1
- "Pinning field columns to help analyze events" on page 52

Exporting search results

About

- ArcSight Investigate can export search results to a CSV file, which can be read by such applications as Excel and Numbers. The data appears in data-table format.
- Exported data is based on the fieldset specified for the search. Filtering has no impact on exported data. If you edit the fieldset, this is reflected in the fields available for the events.
- Hidden columns are exported.
- Investigate preserves any highlights when exporting the search results to a CSV application.

Procedure

Location

Left navigation > Search button > Search page > Events

1. Click .
2. Open the downloaded CSV file in the desired CSV application.

See Also

- "Viewing select event data" on page 53
- "Pinning field columns to help analyze events" on page 52

Appendix A: FAQs

Can I pin a field column in order to compare it against other field values?

Answer

In the Event table, Investigate enables you to pin a field column—to make a column horizontally stationary, in order to better compare the column values against those of other columns.

Related Topic

"Pinning field columns to help analyze events" on page 52

Can I export search-results data to an Excel file?

Answer

Yes. ArcSight Investigate enables you to output search results to any CSV file.

Related Topic

"Exporting search results" on page 54

How much search-result data can I view?

Answer

There are a possible of 176 pieces of data that can be returned from a search. This data takes the form of fields in a fieldset. The fieldset determines what search results data is available for creating data visualization charts and what data columns appear in the Events table.

To manage the display of search-result data, you can limit the number of fields in a fieldset.

Related Topic

"Editing a fieldset for search results " on page 36

Can I view the most and least common values for a search-results field?

Answer

To help you filter for data anomalies, ArcSight Investigate enables you to quickly display the most and least common values for a field. This translates into the count and percentage of hits for the value. For example, the **devicevendor** field could have a top value of "bluecoat" with a count of 3,000 hits, which is 30% of 10,000 results.

Related Topic

"Viewing the most and least common values for an event record field" on page 51

Can I use SQL to specify query input?

Answer

No. Support for SQL statements is planned for a future release.

Can I use a SIEM with ArcSight Investigate?

Answer

Currently, ArcSight Investigate only supports the ArcSight Enterprise Security Manager (ESM) SIEM. Support for other SIEMs is planned for a future release.

Related Topic

- "Features and benefits" on page 4
- "How ArcSight Investigate works" on page 1

Can I apply User Behavior Analytics to the Hadoop data lake used by ArcSight Investigate?

Answer

No. Support for HPE Security ArcSight User Behavior Analytics (UBA) and HPE Security ArcSight DNS Malware Analytics (DMA) to use the Hadoop data lake used by Investigate is planned for a future release.

Related Topic

"Features and benefits" on page 4

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 system, click the link above and an email window opens with the following information in the subject line:

Feedback on User's Guide (ArcSight Investigate 1.01)

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 arc-doc@hpe.com.

We appreciate your feedback!

Glossary

A

Active channel

An active channel is a tool that monitors all the activity that ESM processes for your network. An active channel displays a stream of information defined by parameters set in the active channel editor. A channel could stream events, or show the status of some resources. A channel can be further fine-tuned using in-line filters. There are three types of active channels that display different types of data: - Live Channels continuously refreshed live event data - Rules Channels display replay events for testing rules - Resource Channels display the status of certain resources, such as the assets in your network model and open cases

Aggregate data

Data that refers to numerical or non-numerical information that is (1) collected from multiple sources and/or on multiple measures, variables, or individuals and (2) compiled into data summaries or summary reports, typically for the purposes of reporting or statistical analysis—for such purposes as examining trends, making comparisons, or revealing information and insights that would not be observable when data elements are viewed in isolation.

Aggregate function

In database management an aggregate function is a function where the values of multiple rows are grouped together as input on certain criteria to form a single value of more significant meaning or measurement such as a set, a bag or a list. Common aggregate functions include : Average() and Count().

Apache Avro

Avro is a remote procedure call and data serialization framework developed within Apache's Hadoop project. It uses JSON for defining data types and protocols, and serializes data in a compact binary format. Its primary use is in Apache Hadoop, where it can provide both a serialization format for persistent data, and a wire format for communication between Hadoop nodes, and from client programs to the Hadoop services.

Apache Kafka

Apache Kafka is an open-source stream processing platform. Kafka is a distributed publish-subscribe messaging system that is designed to be fast, scalable, and durable. Like many publish-subscribe messaging systems, Kafka maintains feeds of messages in topics. Producers write data to topics and consumers read from topics. Since Kafka is a distributed system, topics are partitioned and replicated across multiple nodes.

Apache Kafka broker

Each node in an Apache Kafka cluster is called a Kafka broker.

Apache Kafka topic

The container with which messages are associated. A consumer of topics pulls messages off of a Kafka topic while producers push messages into a Kafka topic.

ArcSight Command Center (ACC)
The ArcSight Command Center is a web-based user interface that enables you to perform many of the functions found in the ArcSight Console. ArcSight Command Center provides dashboards, several kinds of searches, reports, case management, notifications, and administrative functions for managing active channels, content, users, connectors, storage, archives, search filters, saved searches, peer configuration, and system logs.

ArcSight Enterprise Security Manager (ESM)
A comprehensive software solution—a SIEM that combines traditional security event monitoring with network intelligence, context correlation, anomaly detection, historical analysis tools, and automated remediation. It consolidates and normalizes data from disparate devices across your enterprise network in a centralized view.

ArcSight Event Broker
Event Broker centralizes event processing, helps you to scale your ArcSight environment, and opens up ArcSight events to ArcSight Investigate. It enables you to take advantage of scalable, high-throughput, multi-broker clusters for publishing and subscribing to event data.

ArcSight Logger
Logger is a log management solution that is optimized for extremely high event throughput, efficient long-term storage, and rapid data analysis. Logger receives and stores events; supports search, retrieval, and reporting; and can optionally forward selected events. Logger compresses raw data, but can always retrieve unmodified data on demand for forensics-quality litigation data.

ArcSight SmartConnector
The interface to the objects on your network that generate correlation-relevant event data. After collecting event data for ArcSight Investigate, the connectors normalize the data in two ways: normalizing values (such as severity, priority, and time zone) into a common format, and normalizing the data structure into a common schema. SmartConnectors can then filter and aggregate events to reduce the volume of events sent to the ArcSight Manager. See SmartConnector documentation for complete details.

Area chart
An area chart or area graph displays graphically quantitative data. It is based on the line chart. The area between the axis and line are commonly emphasized with colors, textures and hatchings. Commonly, one compares with an area chart two or more quantities.

B

Bar chart
A bar chart or bar graph is a chart or graph that presents grouped data with rectangular bars with lengths proportional to the values that they represent. The bars can be plotted vertically or horizontally. A vertical bar chart is sometimes called a Line graph.

C

CEF
Common Event Format (CEF) is an extensible, text-based, high-performance format designed to support multiple device types in the simplest manner possible. Various message syntaxes are reduced to one-matching ArcSight

Enterprise Security Manager (ESM) normalization. Specifically, CEF defines a syntax for log records comprised of a standard header and a variable extension, formatted as key-value pairs. This format contains the most relevant event information, making it easy for event consumers to parse and use them. Other standards target a single component of the security infrastructure or are designed for specific applications. These alternatives lack the ability to support today's high-performance, real-time security requirements. For Investigate, there is CEF to Avro conversion for CEF versions 0.1 and 1.0.

CLI

A command-line user interface (CLI), also known as a console user interface, and character user interface (CUI), is a means of interacting with a computer program where the user (or client) issues commands to the program in the form of successive lines of text (command lines). A program which handles the interface is called a command language interpreter or shell.

Cold data

Data that is accessed less frequently by an organization. Cold data is usually stored on lower performing and less expensive storage environments in-house or in the cloud.

Column chart

A column chart is a graphic representation of data. Column charts display vertical bars going across the chart horizontally, with the values axis being displayed on the left side of the chart.

Connector

An integration element to a certain software, device format, appliance, or function through use of the product. An Onboard Connector means software that resides on the HPE ArcSight appliance that communicates with other software data center. A Remote Connector is software that resides on a different computer that communicates with the HPE ArcSight appliance.

Containers as a Service (CaaS)

To deliver the consistent experience for developers and IT ops, teams began using Docker for Containers as a Service (CaaS). Containers as a Service is a model where IT organizations and developers can work together to build, ship and run their applications anywhere. CaaS enables an IT secured and managed application environment consisting of content and infrastructure, from which developers are able build and deploy applications in a self service manner.

Contextual search

A form of optimizing web-based search results based on context provided to Investigate to execute the query. For example, Investigate knows what operators to provide in the search if an IP address is specified. Likewise, if an operator is specified in the search, Investigate knows what other related operators to provide. When entering query input, Investigate can suggest fields, operators, and searches. The technology understands basic search keywords based on security terminology, database content, and user history. The search is based on such criteria as time, IPs, domains, device vendors, ports, protocols, EventCategory, and usernames.

Continuous data

Data that is not restricted to a specific value, but can occupy any value over a continuous range.

CSV

In computing, a comma-separated values (CSV) file stores tabular data (numbers and text) in plain text. Each line of the file is a data record. Each record consists of one or more fields, separated by commas. The use of the comma

as a field separator is the source of the name for this file format.

D

Data lake
A storage repository that holds a vast amount of raw data in its native format until it is needed. While a hierarchical data warehouse stores data in files or folders, a data lake uses a flat architecture to store data.

Data visualization
Data visualization is the graphical display of abstract information for two purposes: (1) Sense-making (also called data analysis) and (2) communication. Important stories live in data and data visualization is a powerful means to discover and understand these stories, and then to present them to others.

Dataset
A collection of related sets of information that is composed of separate elements but can be manipulated as a unit by a computer.

Discrete data
Data that can be numeric, such as the number of apples. It can also be categorical, like red or blue, or male or female, or good or bad.

Docker
Docker containers wrap a piece of software in a complete filesystem that contains everything needed to run: code, runtime, system tools, system libraries – anything that can be installed on a server. This guarantees that the software will always run the same, regardless of its environment. The Docker platform leverages Docker containers to enable IT operations teams and Developer teams to build, ship and run any application, anywhere. Docker containers are based on open standards, enabling containers to run on all major Linux distributions and on Microsoft Windows -- and on top of any infrastructure. Docker creates a common framework for developers and sysadmins to work together on distributed applications.

F

Fieldset
A select group of fields that determine the field information that displays in the search results for each event that matched the search query. Investigate provides a predefined, default fieldset.

Full-text search
Searches on all the tables. If you enter a string you don't know about you just search the entire columns in all the tables.

H

Hadoop
Apache Hadoop is an open source software platform for distributed storage and distributed processing of very large data sets on computer clusters built from commodity hardware. Hadoop services provide for data storage,

data processing, data access, data governance, security, and operations.

Hadoop cluster
A special type of computational cluster designed specifically for storing and analyzing huge amounts of unstructured data in a distributed computing environment.

Hadoop data lake
A data management platform comprising of one or more Hadoop clusters used principally to process and store non-relational data such as log files , Internet clickstream records, sensor data, JSON objects, images and social media posts.

HDFS
The core of Apache Hadoop consists of a storage part, known as Hadoop Distributed File System (HDFS), and a processing part which is a map-reduce programming model. Hadoop splits files into large blocks and distributes them across nodes in a cluster.

Hot data
Data that needs to be accessed frequently. It is typically business-critical information that needs to be accessed quickly and is often used by a company for quick decision making. Hot data usually resides on the fastest storage -- typically flash in hybrid or tiered storage environments.

HPE Security ArcSight DNS Malware Analytics (DMA)
DMA is a scalable, cloud-based threat detector that monitors DNS traffic and rapidly identifies an infected system, enabling immediate remediation in real time. The application can function in a stand-alone configuration as well as in a Security Operations Center (SOC), using HPE - Security ArcSight Enterprise Security Manager (ESM) as the Security Information and Event Management (SIEM) tool.

HPE Security ArcSight User Behavior Analytics (UBA)
HPE Security ArcSight User Behavior Analytics (UBA) enables security analysts to minimize the risk and impact of cyberattacks in real time. Instead of solely focusing on events and log data, HPE ArcSight UBA detects unknown threats through purpose-built security analytics by creating a baseline of normal user and entity behavior and identifying anomalies associated with users and entities as they occur. By aggregating activities and multiple indicators of compromise for users, entities, and their peer groups, HPE ArcSight UBA delivers insight into the highest risk users and entities—even when credentials are legitimate.

HPE Vertica
An advanced SQL database that can address the most demanding Big Data analytics initiatives. It introduces a unified architecture and advanced in-database analytics capabilities that enable users to conduct sophisticated analysis at industry-leading scale and speed, regardless of where their data resides.

I

Integration commands
Integration commands are a set of tools in the ESM Console that make it possible to invoke scripts and utilities from several places in the ArcSight Console, and to provide snap-in views of other applications, such as ArcSight Logger and third-party applications, within the ArcSight Console. This enables you to use the ArcSight Console as a central command hub for all security-related operations. Once integrated, the commands, tools, and applications

can be launched on demand from within the Console, such as from a right-click context menu within an events grid.

IoT

The Internet of things (IoT) is the inter-networking of physical devices, vehicles (also referred to as "connected devices" and "smart devices"), buildings, and other items—embedded with electronics, software, sensors, actuators, and network connectivity that enable these objects to collect and exchange data.

K

Kubernetes

Commonly referred to as "K8s", this is an open-source container cluster manager originally designed by Google and donated to the Cloud Native Computing Foundation. It aims to provide a "platform for automating deployment, scaling, and operations of application containers across clusters of hosts". It usually works with the Docker container tool and coordinates between a wide cluster of hosts running Docker.

L

Line chart

A line chart or line graph is a type of chart which displays information as a series of data points called 'markers' connected by straight line segments. It is a basic type of chart common in many fields.

M

Microservices

Microservices is a specialisation of and implementation approach for service-oriented architectures (SOA) used to build flexible, independently deployable software systems. As with SOA, services in a microservice architecture (MSA) are processes that communicate with each other over a network in order to fulfill a goal. Also, like SOA, these services use technology-agnostic protocols. The microservices approach is a first realization of SOA that followed the introduction of DevOps and is becoming more popular for building continuously deployed systems. In a microservices architecture, services should have a small granularity and the protocols should be lightweight. A central microservices property that appears in multiple definitions is that services should be independently deployable. The benefit of distributing different responsibilities of the system into different smaller services is that it enhances the cohesion and decreases the coupling. This makes it easier to change and add functions and qualities to the system at any time. It also allows the architecture of an individual service to emerge through continuous refactoring, and hence reduces the need for a big up-front design and allows for releasing software early and continuously.

N

Natural-language search

A set of pre-defined operators. Complex search: Two or more terms Separation operators: 1. And 2. Not 3. = 4. OR 5. Connecting to 6. Equals 7. List (src =1.1.1.1, 1.2.4.5) This is an OR. Example: src = 1.1.1.1 or src = 1.2.4.5

NOC

Network Operations Centers (NOCs) are implemented by business organizations, public utilities, universities, and government agencies that oversee complex networking environments that require high availability. NOC personnel are responsible for monitoring one or many networks for certain conditions that may require special attention to avoid degraded service. Organizations may operate more than one NOC, either to manage different networks or to provide geographic redundancy in the event of one site becoming unavailable. In addition to monitoring internal and external networks of related infrastructure, NOCs can monitor social networks to get a head-start on disruptive events.

O

OT

Operational Technology (OT) is hardware and software that detects or causes a change through the direct monitoring and/or control of physical devices, processes and events in the enterprise.

P

Pie chart

A type of graph in which a circle is divided into sectors that each represent a proportion of the whole. A pie chart can be used to show percentages of a whole, and represent percentages at a set point in time.

R

REST

REST (REpresentational State Transfer) is an architectural style, and an approach to communications that is often used in the development of Web services. The REST architectural style describes six constraints: - Uniform Interface - Stateless - Cacheable - Client-Server - Layered System - Code on Demand (optional)

ROS

The Read Optimized Store (ROS) is a highly optimized, read-oriented, disk storage structure, organized by projection. The ROS makes heavy use of compression and indexing. You can use the COPY statement DIRECT and INSERT parameters (with /*+direct*/ hint) to load data directly into the ROS. Note: HPE Vertica allows optional spaces before and after the plus sign in direct hints (between the /* and the +).

RSS

RSS (Rich Site Summary; originally RDF Site Summary; often called Really Simple Syndication) uses a family of standard web feed formats to publish frequently updated information: blog entries, news headlines, audio, video.

Runbook

In a computer system or network, a runbook is a routine compilation of procedures and operations that the system administrator or operator carries out. System administrators in IT departments and NOCs use runbooks as a reference. Runbooks can be in either electronic or in physical book form.

S

Scatter plot chart
A scatter plot (also called a scatter graph, scatter chart, scattergram, or scatter diagram) is a type of plot or mathematical diagram using Cartesian coordinates to display values for typically two variables for a set of data. If the points are color-coded, one additional variable can be displayed.

Security analyst
The primary user of ArcSight Investigate. This user relies on the overall ArcSight log collection and search capabilities for successfully triaging security incidents. Ultimately, security analysts want to get actionable insights from a search.

Security architect
This user is responsible for determining the overall ArcSight deployment and how this product fits into the SIEM architecture of the organization. This includes integration with other systems such as Hadoop which may be used for storing additional log data.

Security engineer
This user is responsible for data sources and determining how security analysts can effectively triage security incidents and security threat.

Security posture
Your overall security plan – the approach your organization takes to security, from planning to implementation. It is comprised of technical and non-technical policies, procedures and controls, that protect you from both internal and external threats.

SIEM
In the field of computer security, Security Information and Event Management (SIEM) software products and services combine Security Information Management (SIM) and Security Event Management (SEM). They provide real-time analysis of security alerts generated by network hardware and applications. HPE Security ArcSight Enterprise Security Manager (ESM) is an example of a SIEM product.

SMTP
SMTP (Simple Mail Transfer Protocol) is a TCP/IP protocol used in sending and receiving e-mail.

SOC
An information security operations center ("ISOC" or "SOC") is a facility where enterprise information systems (websites, applications, databases, data centers and servers, networks, desktops and other endpoints) are monitored, assessed, and defended.

Sparkline
A small graphic designed to give a quick representation of numerical or statistical information within a piece of text, taking the form of a graph without axes.

System admin
This user is responsible for the deployment, administration and day to day operations of the product. They will need the necessary monitoring and administrative controls to ensure that the product is available and functioning with optimal performance for security analysts.

T

Text box widget
From the Dashboard, you can use this widget to create, edit, and delete a text note.

TLS
Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), both frequently referred to as "SSL", are cryptographic protocols that provide communications security over a computer network. In the case of ArcSight Investigate, TLS is used between the user management, search interface, search engine, and Vertica database modules.

Tuple
A tuple is a sequence of immutable Python objects. Tuples are sequences, just like lists. The differences between tuples and lists are, the tuples cannot be changed unlike lists and tuples use parentheses, whereas lists use square brackets. Creating a tuple is as simple as putting different comma-separated values.

V

Vertica
At its core, the HPE Vertica Analytics Platform from Hewlett Packard Enterprise is a column-oriented, relational database system built specifically to handle modern analytic workloads. The platform uses a clustered approach to storing big data, offering high-performance query and analytics functionality.

Z

ZooKeeper
Apache ZooKeeper is a distributed hierarchical key-value store, which is used to provide a distributed configuration service, synchronization service, and naming registry for large distributed systems. ZooKeeper was originally a sub-project of Hadoop. ZooKeeper's architecture supports high availability through redundant services. The clients can thus ask another ZooKeeper leader if the first fails to answer. ZooKeeper nodes store their data in a hierarchical name space, much like a file system or a tree data structure. Clients can read from and write to the nodes and in this way have a shared configuration service. Updates are ordered.