



**Getting started with
the Siren Platform™**

Table of contents

About this tutorial	4
What does this tutorial cover?	4
What tasks will you complete?	5
What will you get from this tutorial?	5
Installing Siren Platform - Easy Start	6
Before you begin	6
Installing and running Siren Investigate	6
Relaunching Siren Investigate	8
Welcome to Siren Investigate!	8
Importing data into Siren Platform	9
About the demo data	9
Importing your first table of demo data	10
Index pattern searches	15
Creating your first index pattern search	15
Creating dashboards automatically	17
Automatically generating a dashboard	18
Editing dashboards	21
Using the dashboard	23
Creating filters	23
Textual Search	24
Creating dashboards manually	26
Importing the investments table	26
Creating the dashboard	27
Creating visualizations	28
Vertical Bar Chart	28
Record Table	30
Analytic Table	31
Setting the Dashboard Data Model	32
Importing the investors table	33
Introducing relations in the Data Model	35
Creating relations between entities	36
Creating relations manually	37
Visualizing the data model as a graph	39

Creating a self-relation	40
Dashboard-to-dashboard associative navigation	42
Recap: Where are we now?	44
EXERCISE: Investigating investments	46
Part 1: Finding an answer by filtering the investors dashboard	46
Part 2: Finding an answer by using relational navigation	52
Part 3: Digging deeper with relational navigation	55
Part 4: Link analysis in the Graph Browser	57
More about link analysis	66
Embedded in dashboards	66
Using the Graph Browser	66
Appendix 1: Natural Language Processing (NLP) and more link analysis	70
Connecting extracted entities in the data model	73
Example Usage of NLP Data in Siren Investigate	73
Appendix 2: Configuring the graph	77
Graph lenses: Fixing labels with lenses	77
Using graph cards	78
Appendix 3: The Topic Explorer	79
Appendix 4: Creating Entity Identifiers	81
Appendix 5: Dashboard 360	82
Filtering with Dashboard 360	85
Appendix 6: Deleting data in Siren and changing a table schema	87
Appendix 7: Auto-discovery of relations	89
Legal notices	93

About this tutorial

This tutorial will help you to get started with the **Siren Platform - Easy start**, version 11.0 and later, which can be downloaded [here](#).

The tutorial helps you to install Siren Investigate as an empty platform, to import demo data, and to start exploring and analyzing the data.

What does this tutorial cover?

Siren Platform can be applied to many kinds of scenarios and the method of importing data can vary, based on the type of data that is used.

This tutorial focuses on using *static* data, which we will import from Excel or CSV files. Static data, which does not change or changes only a few times per day, is commonly used in Business Intelligence (BI) or broader knowledge-discovery scenarios.

Using streamed data and data from other databases

For use cases that involve streaming data, such as in **cybersecurity**, **operational intelligence**, **log management** or **IoT**, data streams can be loaded by external applications that write directly to the underlying Elasticsearch cluster. Examples of external tools for these scenarios are [Logstash](#), [Beats](#), [Fluentd](#), or [Streamset](#). It is also quite easy to write data to Elasticsearch directly by using the APIs.

Siren Platform also supports working on data in remote JDBC data sources. This can happen directly (with no data copy) or by using the built-in, UI-assisted *reflection* process, where data is copied and periodically refreshed.

These processes are outside of the scope of this tutorial. However, we recommend that you start with this tutorial (which uses CSV files) and then refer to our [documentation](#) to understand how to use data streams and JDBC data sources.

What tasks will you complete?

In this tutorial, we will guide you through the following steps:

1. Installing Siren Platform.
2. Learning how to import a CSV file: Importing the first table (**companies.csv**).
3. Creating an index pattern search.
4. Creating the first dashboard by using the *auto create wizard*.
5. Editing a dashboard by using interactive visual filters and textual search.
6. Importing the next table (**investments.csv**) and creating a dashboard *manually*.
7. Importing the last table (**investors.csv**).
8. Creating an associative data model: Associating **companies** <-> **investments** <-> **investors**.
9. Creating a dashboard manually and including a relational navigator visualization that will allow for dashboard-to-dashboard navigation across all your data.
10. Introduction to the Graph Browser and Link Analysis, including exercises for you to complete.

The tutorial also includes a number of appendixes that contain additional important topics.

What will you get from this tutorial?

By the end of this tutorial, and after completing the steps in Appendix 1, you will have a setup that is similar to our 'What is Investigative Intelligence' video, which we highly recommend you watch now.

Watch the [What is Investigative Intelligence](#) video.

This tutorial also has an accompanying video, which is available on the [Siren website](#).

Installing Siren Platform - Easy Start

This download is supported on Windows, Linux, and macOS operating systems with a minimum of 8 GB of RAM.

Siren Platform supports the following browsers:

- Google Chrome
- Mozilla Firefox

Before you begin

The following are **mandatory requirements** that must be completed before you install the Siren Platform:

- Ensure that **Java JDK version 8** is installed and set the `java_home` environment variable.

NOTE: We recommend downloading Amazon Corretto version 8 (go to the [Windows installer](#) or general [download page](#)). Alternatively, download and install the [Oracle JDK](#). To set the `java_home` environment variable, follow the instructions [here](#).

- On Windows, you must have an **improved ZIP manager** such as [7-zip](#) or [WinRAR](#). This is because the standard windows ZIP support will not be able to handle the distribution's large number of files.

Installing and running Siren Investigate

To install Siren Investigate, complete the following steps:

1. Go to <http://www.siren.io/downloads> and download Siren Platform *Easy start*.
2. Complete the validation form, accept the license, and click **Proceed**.
3. Save the compressed file and extract it to a local directory.

NOTE: On macOS, if you are running Catalina or a later version, you must remove an extended file attribute (xattr) before you extract the downloaded ZIP file. Run the following command:

```
xattr -d com.apple.quarantine
siren-platform-*-darwin-x86_64.zip
```

The extracted folder contains three main folders:

- **docs:** Contains the Siren Platform user guide in both HTML and PDF formats and the release notes.
- **elasticsearch:** Contains the self-contained, single-node Siren Enhanced Elasticsearch cluster.
- **siren-investigate:** Contains the Siren front-end application, which can be accessed through a web browser.

4. Launch the program, based on your operating system.

Windows	<p>Double-click on the start.bat file and wait for Elasticsearch and Siren Investigate to load.</p> <div data-bbox="582 909 1394 1133" style="border: 1px solid blue; padding: 5px;"> <p>NOTE: You will see two command windows; one for Elasticsearch and the other for Siren Investigate. Do not close these windows until you are finished working with Siren Investigate.</p> </div> <div data-bbox="582 1189 1394 1469" style="border: 1px solid red; padding: 5px;"> <p>IMPORTANT: The first time you run the installation, the Investigate window may appear to be stuck for a number of minutes and appear black. This is the optimization phase. Please wait and it will eventually start writing log messages again.</p> </div>
Linux or macOS	<ol style="list-style-type: none"> 1. Open a Terminal window and change directory to the extracted folder: <pre>cd {extracted folder}/elasticsearch</pre> Run the following command: <pre>./bin/elasticsearch</pre> The Terminal window shows Elasticsearch messages flowing. Do not close the command window.

	<p>2. Open a new Terminal window and change directory to the extracted folder:</p> <pre>cd {unzipped folder}/siren-investigate</pre> <p>Run the following command:</p> <pre>./bin/investigate</pre>
<p>NOTE: The installation is complete when you see the message “Siren Gremlin Server is up and running” in the log window.</p>	

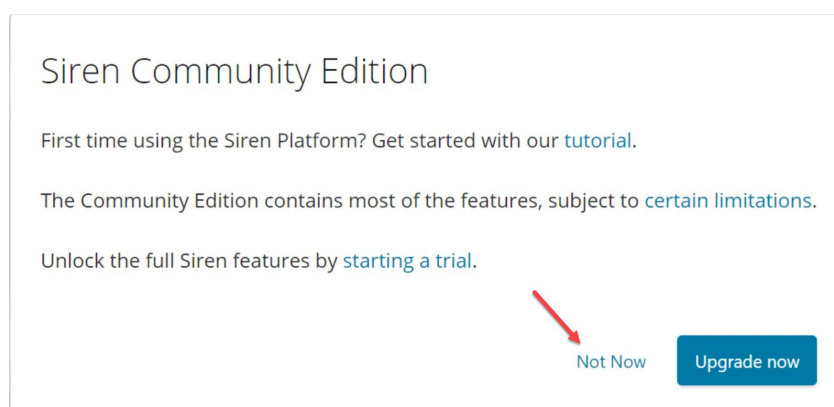
5. Open a web browser and navigate to <http://localhost:5606>.

Relaunching Siren Investigate

If your session is interrupted and you need to relaunch Siren Investigate, repeat steps 4-5 of the above procedure.

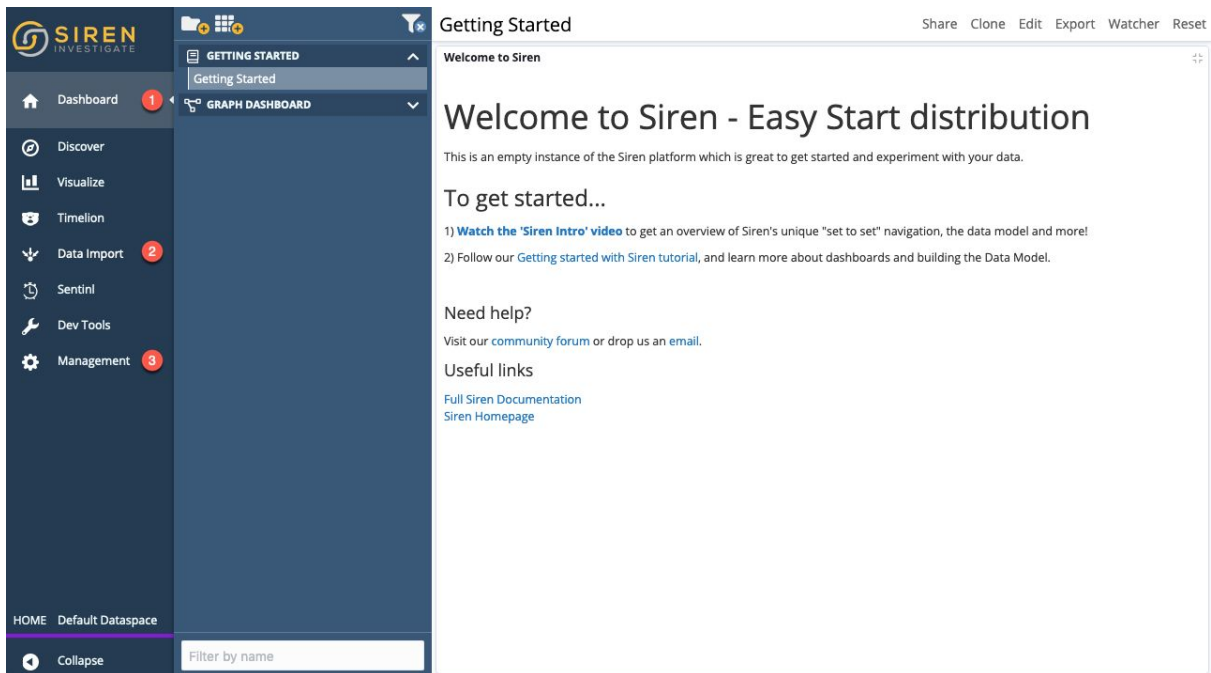
Welcome to Siren Investigate!

As you launch Siren Investigate, you will see our welcome message with some useful links. Click **Not now** to close it.



The platform opens to a welcome dashboard, as you can see in the following screenshot. The left-hand navigation menu contains a range of options. Expand the sidebar to see the most important apps for this tutorial, which are:

- Dashboard ①
- Data Import ②
- Management ③



Importing data into Siren Platform

(Average completion time: 10 mins)

The first step is to import some data that we can work with.

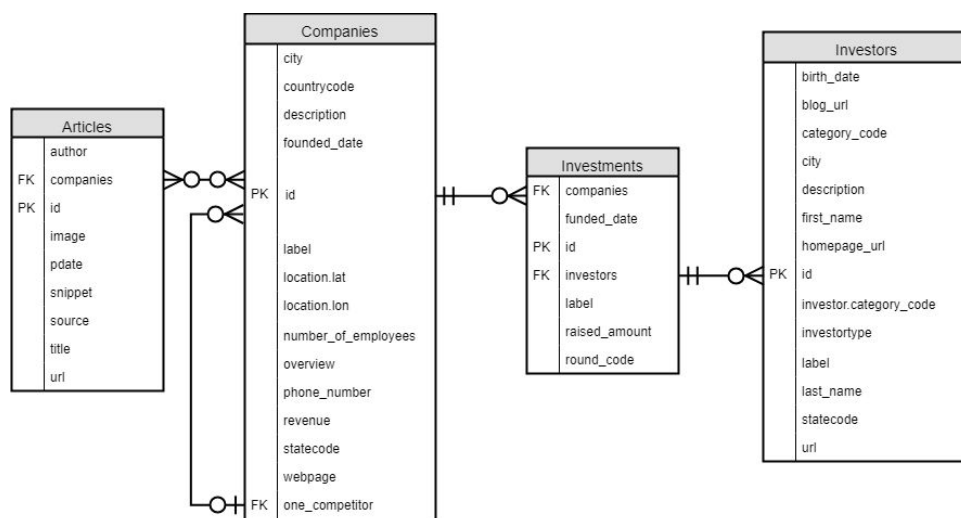
About the demo data

For the purpose of this tutorial, we are using sample data from a TechCrunch data source collected some years ago, along with a sample of technical articles that were collected online.

1. Download the [sample data file](#).
2. Extract the compressed folder that contains the following files:
 - **companies.csv:** A list of companies that includes geo-locations and descriptions.

- **investments.csv:** An associative table that connects companies to investors with “amounts”, “round code” (for example, seed round, or round A), and the date of the investment.
- **investors.csv:** A list of investors.
- **articles.csv:** A collection of technical articles. **Advanced:** Most of the articles mention one or more companies. To extract (or annotate) these mentions, use the Siren NLP capabilities as described in [Appendix 1](#).

These files are relationally connected. The following entity-relationship diagram shows the connections between the tables:



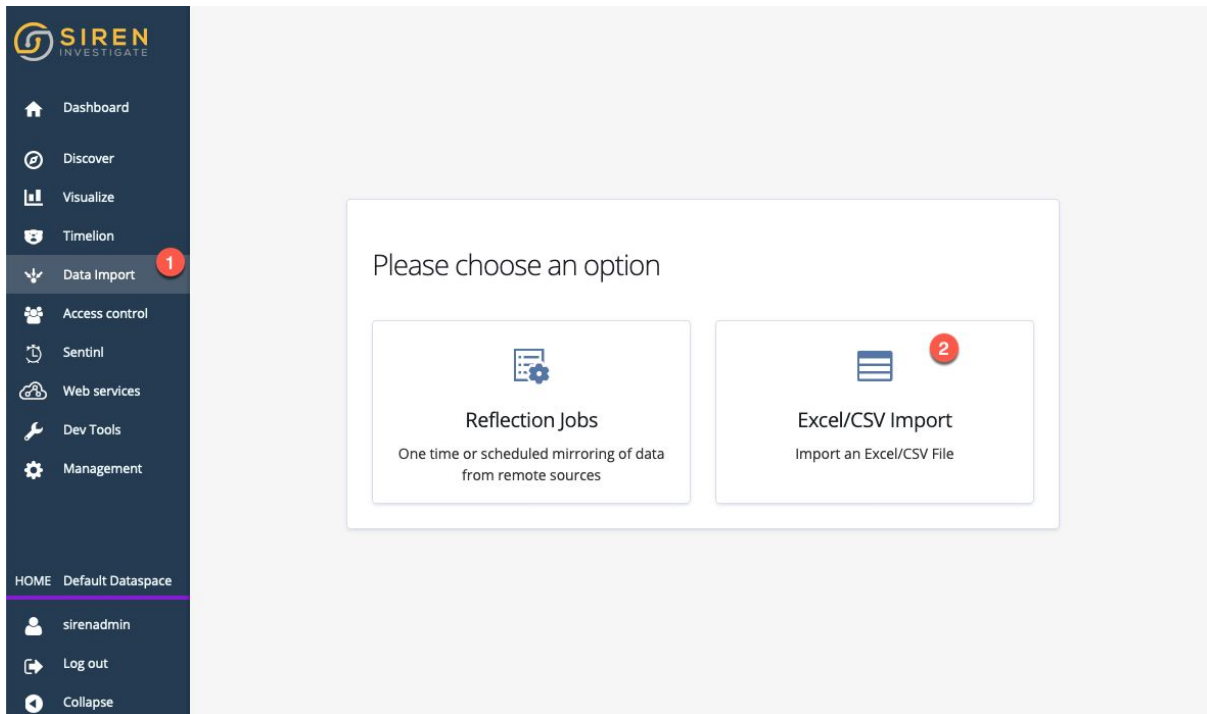
There are also other relationships in the data, for example, city and state names, URLs, and email addresses. You'll see later how these can be auto-discovered.

Importing your first table of demo data

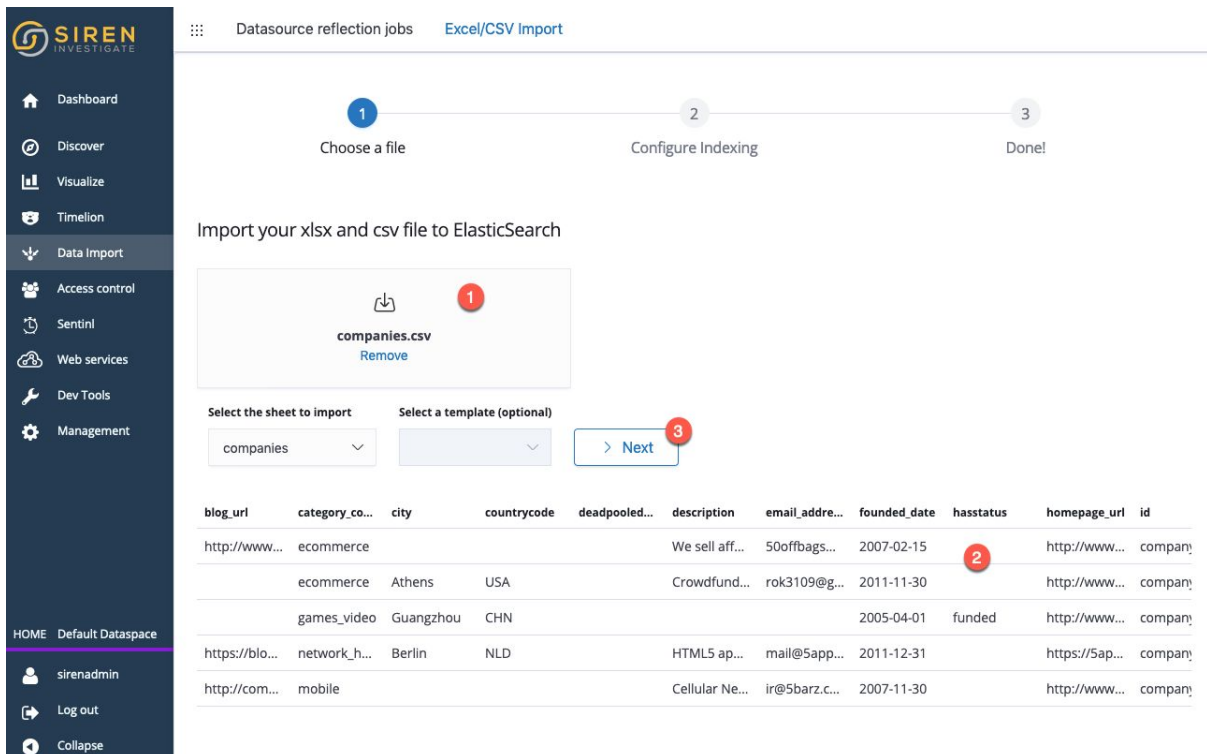
(Average completion time: 5 mins)

This section will take you through the steps required to import your first table of the sample dataset. Let's start by importing the **Companies.csv** file.

1. From the left-hand navigation bar, go to **Data Import** ①.
2. Click on **Excel/CSV Import** ②.



3. From the **Excel/CSV Import** screen, upload the **companies.csv** file ①.
4. You will see a preview of the uploaded data ②. Click **Next** ③.



You will now be on the **Configure Indexing** screen. This screen allows you to make a number of configurations, such as:

- Defining a transform pipeline ①. Accept the default setting for now (**Skip transform**). This is used later to activate transformations such as the [NLP textual processing](#).
- Selecting the **Index scope** ②, which defines the creation of dataspace. Select **Local** from the dropdown menu. For more information, see the [dataspace documentation](#).
- Changing the name of the index ③. Accept the default name for now.
- The **Use Custom Mapping** switch ④ allows you to map the type (or format) for each field in the CSV file. This is very important and determines your ability to search/sort/drill down and analyze the data. Switch **Use Custom Mapping** 'on', because we do need to set the type for a few imported fields.

Most of the data types that are listed in the **Type** drop-down menu ⑤ are easy to understand, such as Integer, or Date. You can pick the correct value by looking at the **Samples** column. However, *text* and *keywords* require some explanation:

Text: Use this option for long texts, such as emails and messages. When you choose **Text**, the back-end splits the content into individual words and calculates statistics on those words. This allows you to have word clouds and the best level of searchability.

Keywords: Use this option for strings that should be considered *unbreakable*. For example, your dataset might include a field called City that contains strings such as 'New York' and 'Los Angeles'. If you mark this field as **Keyword**, then the term will be searched for in its entirety, rather than broken into single-word search terms (which would then lead to wrong results later).

The **Multi-Valued** switch ⑥ can be used when a field in the CSV file contains an array of values, typically separated by a delimiter. For example, you might have an **email** field that contains "mario@bros.io|mario.bros@mail.io|info@mario.io". By switching the **Multi-Valued** switch to 'on', you can parse these three values as separate email addresses, which is useful for search and analysis.

The **Advanced Mapping** switch ⑦ gives advanced users the option to add a JSON object to define the mapping in Elasticsearch format. For example, selecting **Date** and switching

the **Advanced Mapping** switch to 'on' would allow you to define a specific format for the date field, such as {"format": "yyyy-MM-dd HH:mm"}.

Choose a file 2 3 Done!

Transform Pipeline

Method: Skip transform (1)

Indexing Settings

Index scope: Local (2)

Index name: siren-import-home-companies (3)

☐ Use Custom storage primary key?

Remove columns:

Mapping

☒ Use Custom Mapping? (4)
Define Index Mapping.

Fields	Samples	Type	Remove Column
blog_url	blog_url Samples http://www.50offbags.com/blog/ https://blog.5apps.com http://community.5barz.com/	Keyword (For Aggregations) (5) <input type="radio"/> Multi-Valued? (6)	<input checked="" type="checkbox"/> Advanced Mapping? (7)

For now, complete the following steps for the **companies.csv** file:

1. Scroll to the **description** field and set the type to **Text (Allows Word Cloud)**.
2. For the **founded_date** field, set the type to **Date**.
3. For the **Geopoint** field, set the type to **Geo Point**.
4. For the **overview** field, set the type to **Text (Allows Word Cloud)**.
3. Click **Import** ①.

SIREN INVESTIGATE

- Dashboard
- Discover
- Visualize
- Timelion
- Data Import**
- Access control
- Sentinel
- Web services
- Dev Tools
- Management

HOME Default Dataspace

sirenadmin
Log out
Collapse

phone_number Samples

unknown
5132050945
949-916-3261

phone_number

Keyword (For Aggrega...
Multi-Valued?

statecode Samples

OH

statecode

Keyword (For Aggrega...
Multi-Valued?

url Samples

http://www.crunchbase.com/company/50offbags
http://www.crunchbase.com/company/530funds
http://www.crunchbase.com/company/56-com
http://www.crunchbase.com/company/5apps
http://www.crunchbase.com/company/5barz-international

url

Keyword (For Aggrega...
Multi-Valued?

< back ☒ Auto-Stop at: 150 Errors **Import** **Save template**

After the import is complete, you see a confirmation message.

The data has now been imported (or indexed) in the back-end system. Next, we need to access it through the **Data model** in the front-end system.

To do this, click **Create an Index Pattern Search** ① in the confirmation message.

⋮ Datasource reflection jobs [Excel/CSV Import](#)

Choose a file Configure Indexing Done!

✓ Your file has been imported!

160026 out of 160026 document(s) have been imported into **companies**.

To use this data, you may want to **Create an Index Pattern Search** ① for this Index.

File name : companies.csv

Sheet name : companies

Ingestion Rate: 3355.26 Docs/Sec

Time Taken: 00:00:47.694

TIP: You can also create an index pattern by going to **Management → Data Model → Create index pattern search** ①.

Management / Datamodel / New / Search Entity / Siren Import Home Companies

Data Model | Saved Objects | Templates | Scripts | License | Data Sources | Virtual Indices | Icon Packs | Advanced Settings

+ Create index pattern search ①

+ Create entity identifier

There are no Index Pattern Searches or Entity Identifiers

Create an Index Pattern Search

Siren Investigate requires at least one configured Index Pattern Search. Index Pattern Searches identify the indexes to run.

Index pattern

siren-import-home-companies

The Index Pattern selects the indexes to query.
You can set the name of the index or select multiple indexes with the * wildcard, for example: logstash-*.
Note: virtual indexes can be used only by direct naming, not as part of wildcard patterns.

Time Filter field name refresh fields

Name of this search

siren-import-home-companies

The search name, typically describing returned entity types, for example: "users", "logs firewall 1", "articles.

Short Description

Icon

fas fa-circle

Color

☒ Generate the *primary key* and *single value* field properties automatically

Time Filter field name is required

HOME

Index pattern searches

Before we create an index pattern search, let's present some fundamental concepts and terminology that will help you to understand the platform.

- **A document:** A single record. In this case, we have created one document for each row in the CSV file. Documents have *fields* that contain *values*.
- **An index:** In Elasticsearch, an index is a concept very similar to a table in a database (a collection of documents). The CSV import has created the **companies** index for us at this point.
- **An index pattern search:** A way to combine more indexes. The simplest form of an index pattern search is... simply the name of an index itself, for example, *companies*. More advanced index patterns can be used to combine indexes in advance deployments, for example, indexes that are suffixed or prefixed with dates.

Index pattern searches are the core ‘searchable objects’ in Siren Platform. This means that index patterns can be searched, drilled down, and become the basis of analytical visualizations.

- **Searches or ‘saved searches’:** Typically, searches are a subset of an index pattern search. For example, if you wanted to sort your list of companies by those based in France, you can set up a search for French companies, which would then appear under the ‘companies’ index pattern search.

Creating your first index pattern search

(Average completion time: 5 mins)

To create your first index pattern search, complete the following steps:

1. Keep the default name of the index pattern ① and keep the name of this search ③ as “companies”.
2. In the **Time Filter field name** field ②, select **founded_date** from the drop-down menu.
TIP: If you choose a Time Filter field name, then Siren Platform treats the data as a *Time Series*. In this case, we will consider companies to be a time series and regard the “founded_date” as the main time¹.
3. Select an icon ④ to represent the type of data that the index pattern will show. For example, an icon of a building could represent a company.
4. Click **Create** ⑤.

¹ In general, if you select a time field, then this must be present in all of your records or you will not see all of your data.

Management / Datamodel / New / Search Entity / Siren Import Home Companies

Data Model Saved Objects Templates Scripts License Data Sources Virtual indices Icon Packs Advanced Settings

+ Create index pattern search

+ Create entity identifier

There are no Index Pattern Searches or Entity Identifiers

Create an Index Pattern Search

Siren Investigate requires at least one configured Index Pattern Search. Index Pattern Searches identify the indexes to run.

Index pattern

siren-import-home-companies 1

The Index Pattern selects the indexes to query.
You can set the name of the index or select multiple indexes with the * wildcard, for example: logstash-*.
Note: virtual indexes can be used only by direct naming, not as part of wildcard patterns.

Time Filter field name ? refresh fields

founded_date 2

Name of this search

siren-import-home-companies 3

The search name, typically describing returned entity types, for example: "users", "logs firewall 1", "articles."

Short Description

Icon

far fa-building 4

Color

#ad1ac4

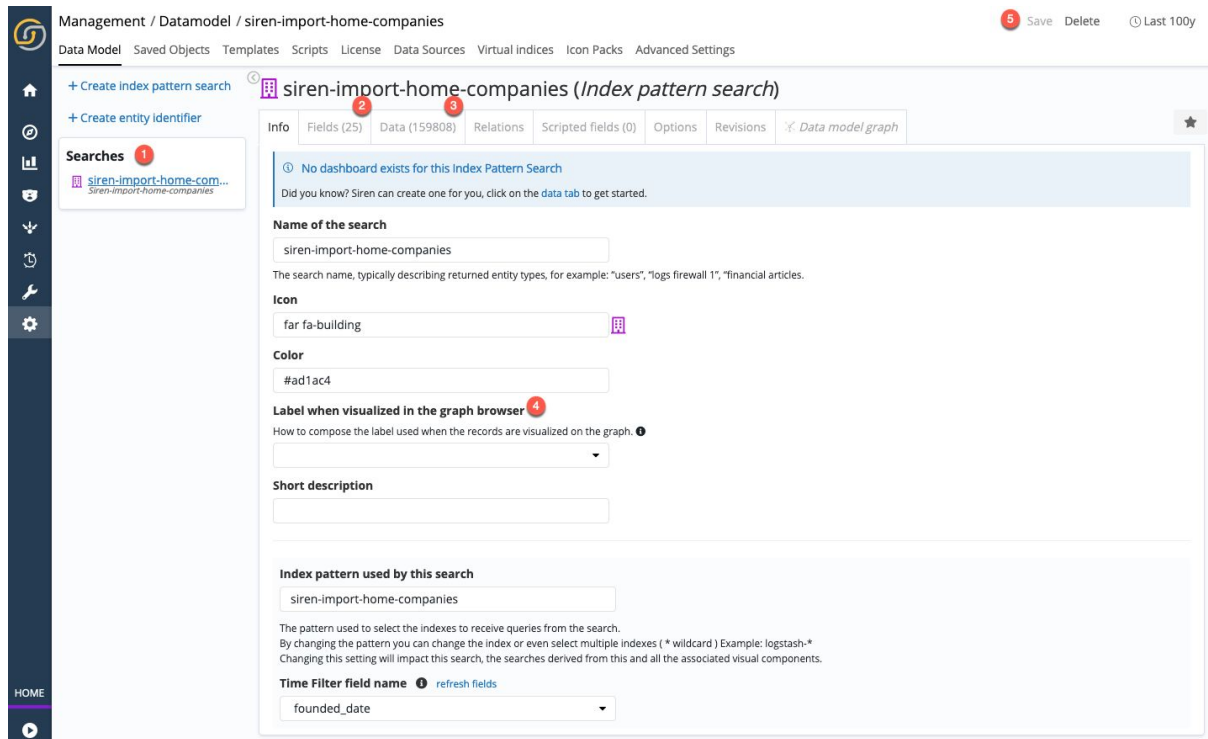
☒ Generate the *primary key* and *single value* field properties automatically ?

Create 5

HOME

After you've created the **Index Pattern Search** you should see something like the screenshot below.

All index pattern searches that you create appear on the left under **Searches** ①.



5. Select the **Fields** tab ② to see which fields are now imported.
6. Select the **Data** tab ③ to see the imported records: Approximately 160,000 records are imported for this dataset.
7. On the **Info** tab, the field **Label when visualized in the graph browser** ④ is used to set the default field for labeling nodes on the graph browser. Select **Document Field** and, in the second dropdown menu that appears, select **label**.
8. Click **Save** ⑤.

Creating dashboards automatically

The Siren Investigate dashboards display a set of visualizations in a customizable grid layout. Each visualization is normally linked to a search (*such as companies or investments*). Dashboards can be organized in dashboard groups, shared, and customized.

You can create dashboards manually or automatically. To manually create a dashboard, you'd need to create and add each visualization manually to your dashboard. Alternatively, the dashboard generator automatically creates the visualizations for you.

The following message appears after you create a Index Pattern Search. Let's see together how we can automatically create a dashboard with Siren Investigate.

① No dashboard exists for this Index Pattern Search
Did you know? Siren can create one for you, click on the [data tab](#) to get started.

Automatically generating a dashboard

(Average completion time: 5 mins)

Siren features an AI-driven dashboard creation wizard, which is a great way to get started on any new dataset with just a few clicks.

The auto-generate dashboard option is available both on:

- the **Data Model** screen on the **Data** tab and;
- the **Discover** app in the navigation menu, which is a stand-alone app that provides a shortcut to this functionality.

Let's start by generating a dashboard for the **companies** data that we just imported.

1. From the **Data Model** page, go to the **Data** tab ①. This page shows a preview of the records in the index pattern and on the left you can see a list of available fields.
2. Click the **Autoselect Most Relevant** button ②.

Management / Datamodel / siren-import-home-companies

Data Model Saved Objects Templates Scripts License Data Sources Virtual indices Icon Packs Advanced Settings

+ Create index pattern search + Create entity identifier

Searches

siren-import-home-com...
Siren-import-home-companies

Info Fields (25) Data (159808) Relations Scripted fields (0) Options Revisions Data model graph

Selected Fields

? _source

Autoselect Most Relevant

Available Fields

@timestamp 15,000

November 9th 1920, 16:11:06.392 - November 9th 2020, 16:11:06.392

① No dashboard exists for this Index Pattern Search
To auto generate a dashboard press "Autoselect Most Relevant" button on the left or manually select fields to appear.

Search... (e.g. status:200 AND extension:PHP)

Add a filter +

3. The system identifies the top ten fields that are useful in generating the dashboard. You can revise them if you wish and then click **Ok** ①.

Autoselect Top 10 – Report

<input type="checkbox"/>	Type	Name	Chart	Score	Notes
<input type="checkbox"/>	t	blog_url	Analytic Table	0	Less than 1% of documents in 50 terms
<input checked="" type="checkbox"/>	t	category_code	Vertical Bar	37.531	
<input checked="" type="checkbox"/>	t	city	Analytic Table	10.632	Less than 19% of documents in 50 terms
<input checked="" type="checkbox"/>	t	countrycode	Analytic Table	12.234	
<input type="checkbox"/>	t	deadpooled_date	Analytic Table	0	Less than 1% of documents in 50 terms
<input checked="" type="checkbox"/>	t	description	Tag Cloud	36.859	Analyzed
<input type="checkbox"/>	t	email_address	Analytic Table	0	Less than 1% of documents in 50 terms
<input checked="" type="checkbox"/>	t	hasstatus	Pie	34.223	
<input type="checkbox"/>	t	homepage_url	Analytic Table	0	Less than 1% of documents in 50 terms

1 Ok Cancel

4. Click **Generate Dashboard** ②.

Management / Datamodel / siren-import-home-companies

Data Model Saved Objects Templates Scripts License Data Sources Virtual indices Icon Packs Advanced Settings

+ Create index pattern search + Create entity identifier

Searches

siren-import-home-com...
Siren-import-home-companies

siren-import-home-companies (Index pattern search)

Info Fields (25) Data (159808) Relations Scripted fields (0) Options Revisions Data model graph

Selected Fields

- Geopoint
- t category_c...
- t city
- t countrycode
- t description
- t hasstatus
- # number_of...
- t overview
- t statecode

Generate dashboard ②

Available Fields

- @timestamp
- t _id

No dashboard exists for this Index Pattern Search

To auto generate a dashboard press "Autoselect Most Relevant" button on the left or manually select fields.

Search... (e.g. status:200 AND extension:PHP)

Add a filter +

November 9th 1920, 16:11:06.392 - November 9th 2020, 16:11:06.392

Count

15,000

10,000

5,000

0

1930 1940 1950 1960 1970 1980

founded_date per year

Time

August 1st 2014, 01:00:00.000

Geopoint

-

category_code

software

city

-

countrycode

-

description

collat
tion,

5. On the **Create a Prepopulated Dashboard** screen, specify the dashboard title "Companies" and click **Create** ①.

Create a Prepopulated Dashboard (beta) ✕

Title

siren-import-home-companies

☒ **Store time with dashboard**
☒ **Add Related Dashboards visualization**
☐ **Add Multichart visualization (slow)**

Create Associated Saved Search

☒ Save current search "siren-import-home-companies" and use it
☐ Use current search "siren-import-home-companies" without saving
☐ New saved search named "siren-import-home-companies #2"

Create

6. Once the process is complete, the **Generate Dashboard - Report** screen is displayed:

Generate Dashboard – Report

<input checked="" type="checkbox"/>	Type	Field	Chart Type	Vis Name
<input checked="" type="checkbox"/>	🕒	founded_date	Line	Documents count by founded_date
<input checked="" type="checkbox"/>	🌐	Geopoint	Enhanced Coordinate Map	Geopoint locations map
<input checked="" type="checkbox"/>	t	category_code	Vertical Bar	Top 50 category_code
<input checked="" type="checkbox"/>	t	city	Analytic Table	List of city
<input checked="" type="checkbox"/>	t	countrycode	Analytic Table	List of countrycode
<input checked="" type="checkbox"/>	t	description	Tag Cloud	Top 30 description
<input checked="" type="checkbox"/>	t	hasstatus	Pie	Top 30 hasstatus
<input checked="" type="checkbox"/>	#	number_of_employees	Vertical Bar	Histogram of number_of_employees
<input checked="" type="checkbox"/>	t	overview	Tag Cloud	Top 30 overview

Ok **Cancel**

The report screen contains the following elements:

- ① - Checkboxes that allow you to individually select which of the auto-generated visualizations to keep in the new dashboard.

② - Dropdown menus, where you can select the best type of visualization (Chart Type) for each field.

③ - Text boxes that allow you to personalize the name of each visualization (Vis Name).

Click **Ok** ④.

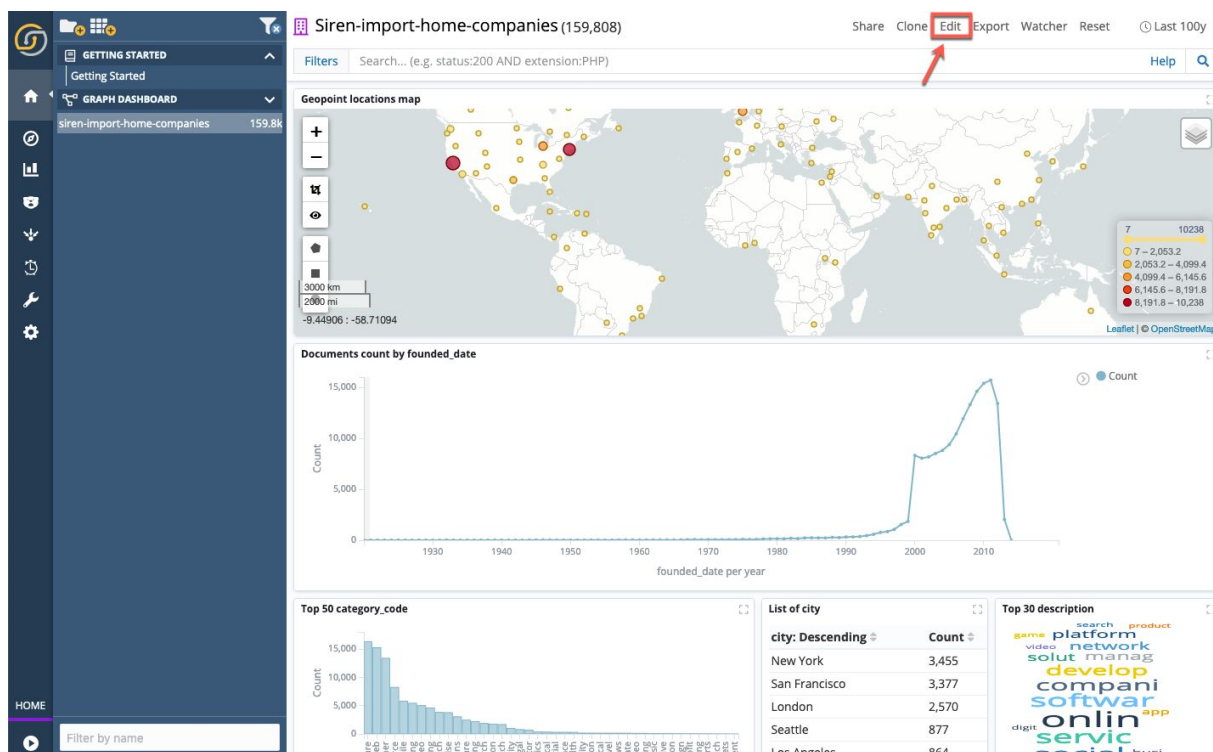
The new **companies** dashboard is now generated with different types of visualizations including an Enhanced Coordinate Map, Histogram, Pie Chart, Record Table, Analytic Table and a Tag Cloud. For more information about visualizations, see the [Siren Platform documentation](#).

TIP: The new dashboard name defaults to the name of the index pattern search, but you can edit the dashboard name at any time by right-clicking the dashboard in the sidebar and clicking **Rename**.

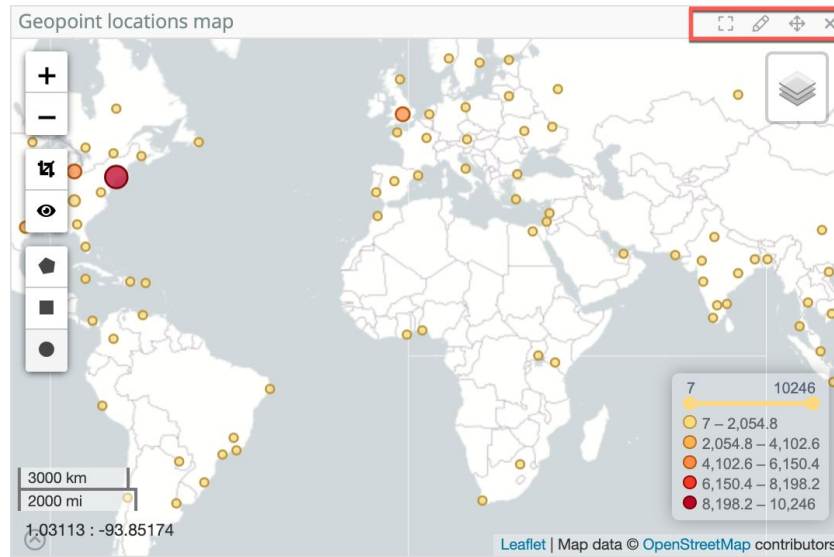
Editing dashboards

(Average completion time: 5 mins)

Let's tidy up the **Companies** dashboard by reorganizing the visualizations.

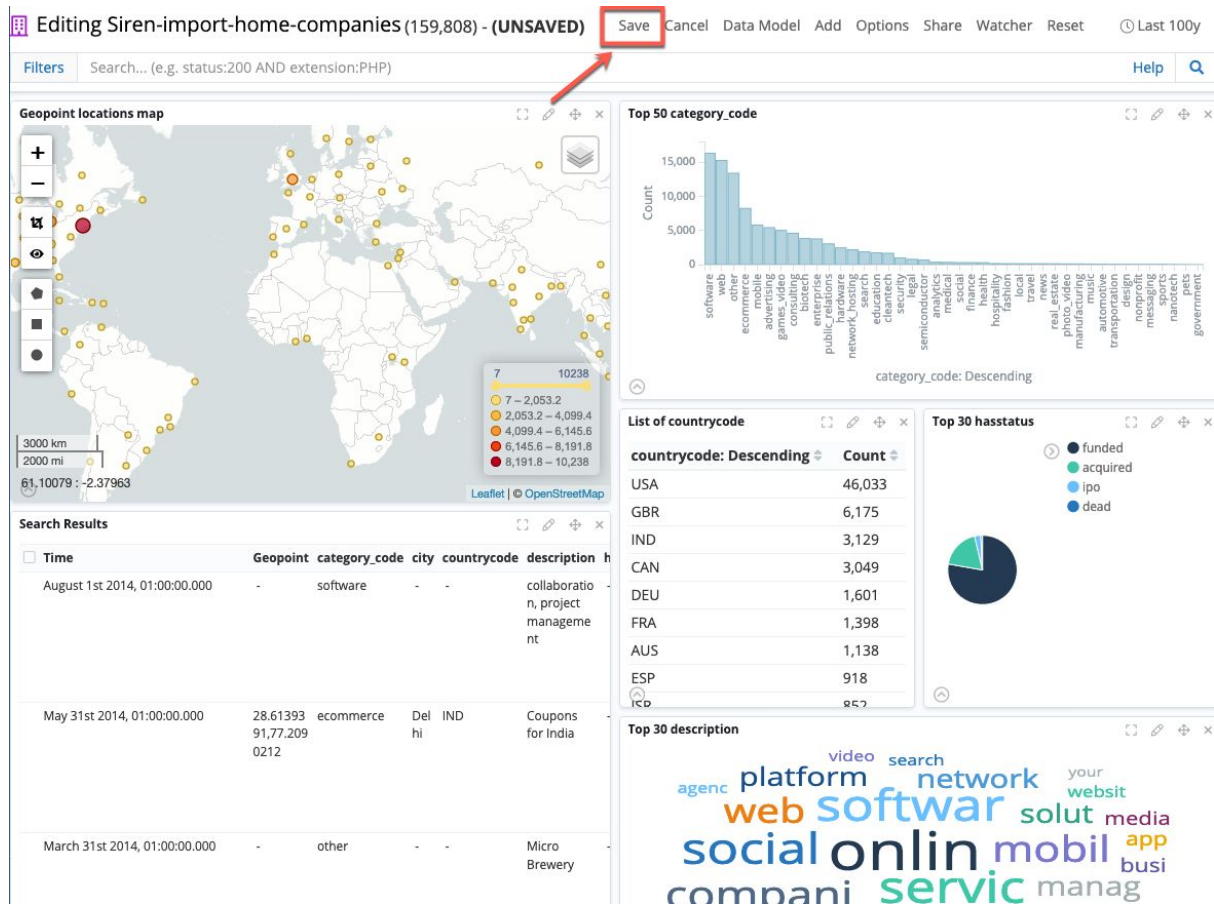


At the top-right of the screen, click **Edit**. The dashboard is now in edit mode and you will notice that the edit buttons appear at the top-right of each visualization. You can now resize, reorder, edit, and remove the visualizations.



Let's rearrange some visualizations to improve the look and feel of the dashboard. You can also remove the visualizations that you think are not relevant.

See below an example of how you can rearrange the companies dashboard:



Click **Save** at the top-right of the screen to save the changes.

Using the dashboard


Now, let's explore this dashboard by interacting with the visualizations to take a closer look at important data.

Creating filters

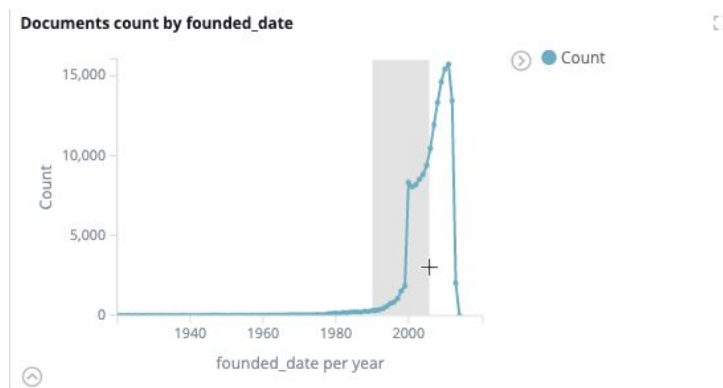
(Average completion time: 5 mins)

You can create filters for dashboards either by using the **Add filter** button or from the visualizations, for example, as follows:

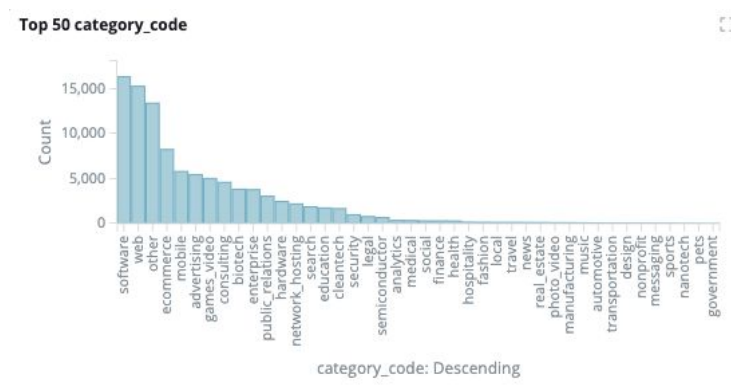
1. Click a segment of the **Top 30 hasstatus** pie chart to filter the dashboard on a category, such as the **funded** status.



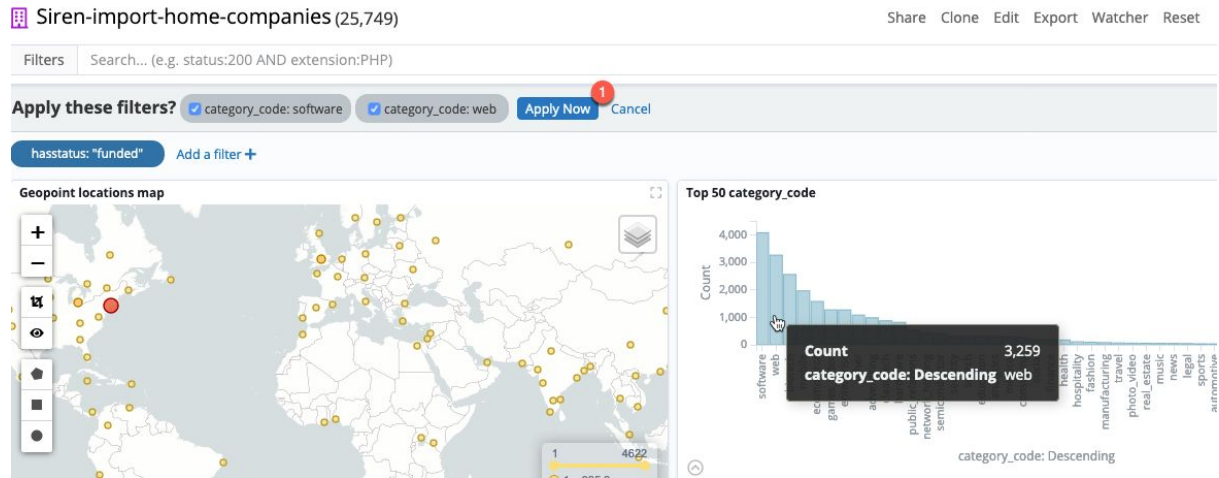
2. On histogram visualizations, you can create multiple types of filters:
- On the **Documents count by founded_date** histogram, click and drag across multiple points to select a range (for time-based histograms only). For this tutorial, there is no need to select a range here right now.



- On the **Top 50 category_code** histogram, click on a bar in the chart to select a specific category.



- You can also make multiple selections by holding **Ctrl** (or **Cmd**) and clicking on the bars. Hold the **Ctrl** key and click on the bar corresponding to **software** and **web** category_code. After making the selections, click on **Apply Now** ① to apply a filter with an OR condition.



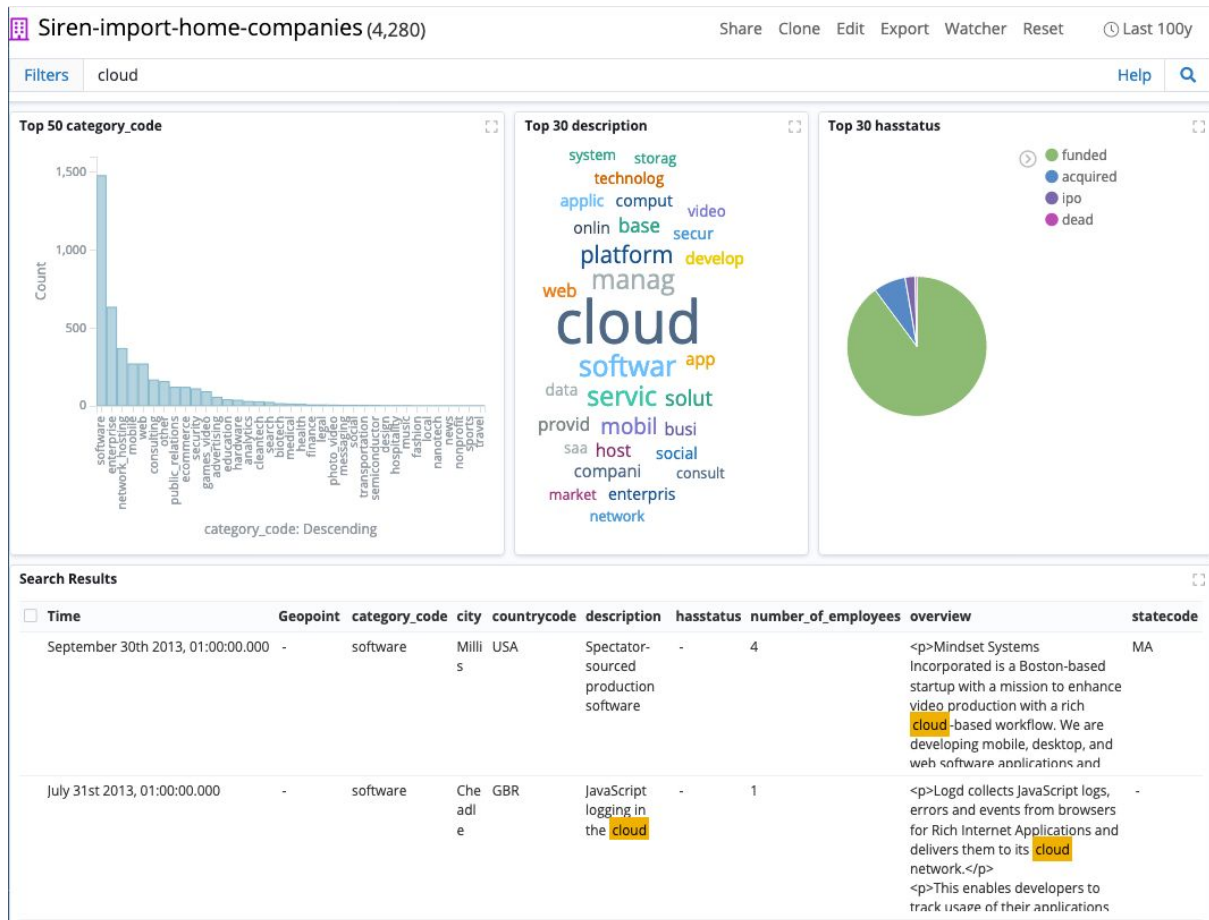
With the filters created, the dashboard updates as follows:

- The following filters appear: **category_code: software,web** and **hasstatus: "funded"**.
- The number of records displayed in the dashboard is reduced and a filter icon appears alongside the dashboard in the sidebar.

Textual Search

On the dashboards, you can perform a free text search by using the Search bar. You can rearrange the visualizations to bring the text-specific ones to the top.

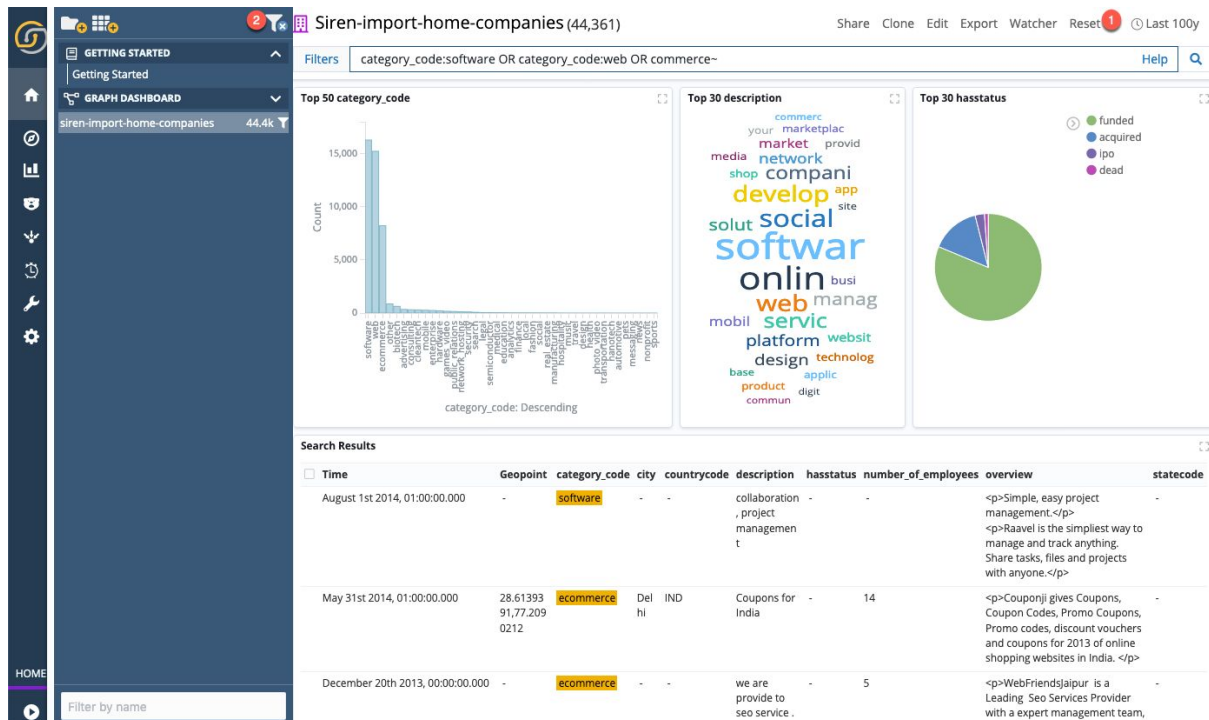
Clear the current filters and simply enter a text string, for example, enter the word "cloud" and hit the **Enter** key. The dashboard updates all of the visualizations to reflect the search results.



You can also perform an advanced search by using the Lucene query syntax.

- To search for a value in a specific field, prefix the value with the name of the field. For example, enter **category_code:software** to find all of the records that have “software” as the category code.
- Use the Boolean operators AND, OR, and NOT for complex queries. For example, enter **category_code:software OR category_code:web** to get all of the results with either the “software” or “web” category code.
- You can also use the **fuzzy operator** ‘~’ to search for terms that are similar to, but not exactly matching the search terms. For example, type **commerce~** to get the results of all mentions of ‘ecommerce’, ‘commercial’, and so on.

To reset the filters on the current dashboard, click **Reset** ①. If you have multiple dashboards, you can reset all of the unsaved filters and searches in all of your dashboards, by clicking **Reset the time, filters, and queries from all dashboards to their default saved state** ②.



Creating dashboards manually

It's time to import your next table and create the next dashboard.

Importing the investments table

(Average completion time: 5 mins)

Return to the **Data Import** app to import the **investments.csv** file and create the relative index pattern search: **Do not** automatically create a dashboard this time.

During the import, make the following settings:

1. Set the **Index scope** to **Local**.
2. For the **funded_date** field, select **Date**.

- For the **raised_amount** field, select **Long**. (This allows an integer over 32 bits: Some investments are big!)
- Click **Import -> Create an Index Pattern Search**.
- On the **Create an Index Pattern Search** screen:
 - In the **Time Filter field name** field, select **funded_date**.
 - Choose an appropriate icon for investments and click **Create**.
- On the **siren-import-home-investments (Index pattern search)** screen, set the **Label when visualized in the graph browser** to be a **Document Field** and, in the second dropdown menu that appears, select the field **raised_amount**.
- Click **Save** in the top-right corner of the screen.

At this point, the **Data Model** tab should look like this:

The screenshot shows the Siren Platform interface for configuring an Index Pattern Search. The breadcrumb trail is 'Management / Datamodel / siren-import-home-investments'. The top navigation bar includes 'Data Model', 'Saved Objects', 'Templates', 'Scripts', 'License', 'Data Sources', 'Virtual indices', 'Icon Packs', and 'Advanced Settings'. The main title is 'siren-import-home-investments (Index pattern search)'. Below the title are tabs for 'Info', 'Fields (19)', 'Data (41317)', 'Relations', 'Scripted fields (0)', 'Options', 'Revisions', and 'Data model graph'. A message states: 'No dashboard exists for this Index Pattern Search. Did you know? Siren can create one for you, click on the data tab to get started.' The configuration fields are:

- Name of the search:** siren-import-home-investments
- Icon:** fal fa-money-bill-alt
- Color:** #27c5f5
- Label when visualized in the graph browser:** Document Field (dropdown), raised_amount (dropdown), with a preview button.
- Short description:** (empty text field)
- Index pattern used by this search:** siren-import-home-investments
- Time Filter field name:** funded_date (dropdown), with a refresh fields link.

Creating the dashboard

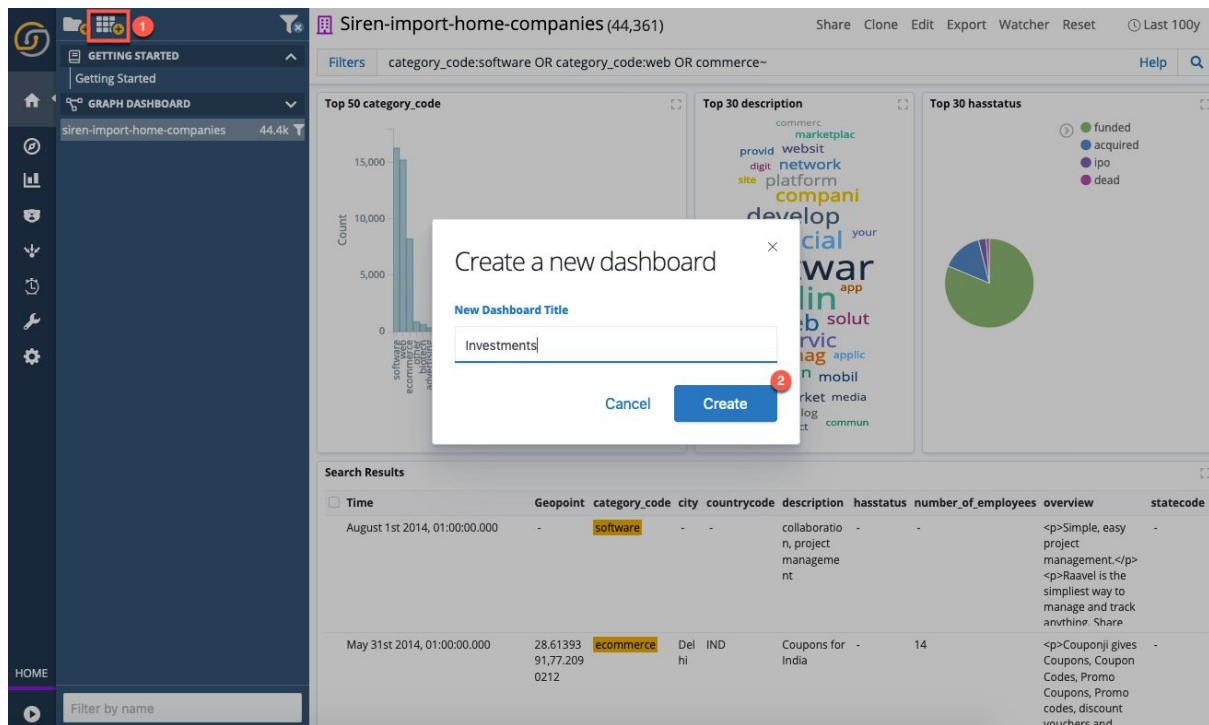
(Average completion time: 7 mins)

In this section, we'll create a dashboard manually for the data that we just imported.

Note: This is an optional step, if you don't want to learn how to do this, you can use the *auto-generate* option as before and skip to the [Importing the investors table](#) section.

To create a dashboard manually for **investments**:

1. Go to the **Dashboard** app.
2. Click on the **Create new dashboard** button ①.



3. Give the dashboard a name, for example, "Investments" and click **Create** ②.

At this point, you have created a dashboard, but it does not yet contain any visualizations and it is not linked to a search.

Creating visualizations

(Average completion time: 15 mins)

You can create visualizations in two ways; either from the **Visualize** tab on the left navigation bar or directly from the current dashboard.

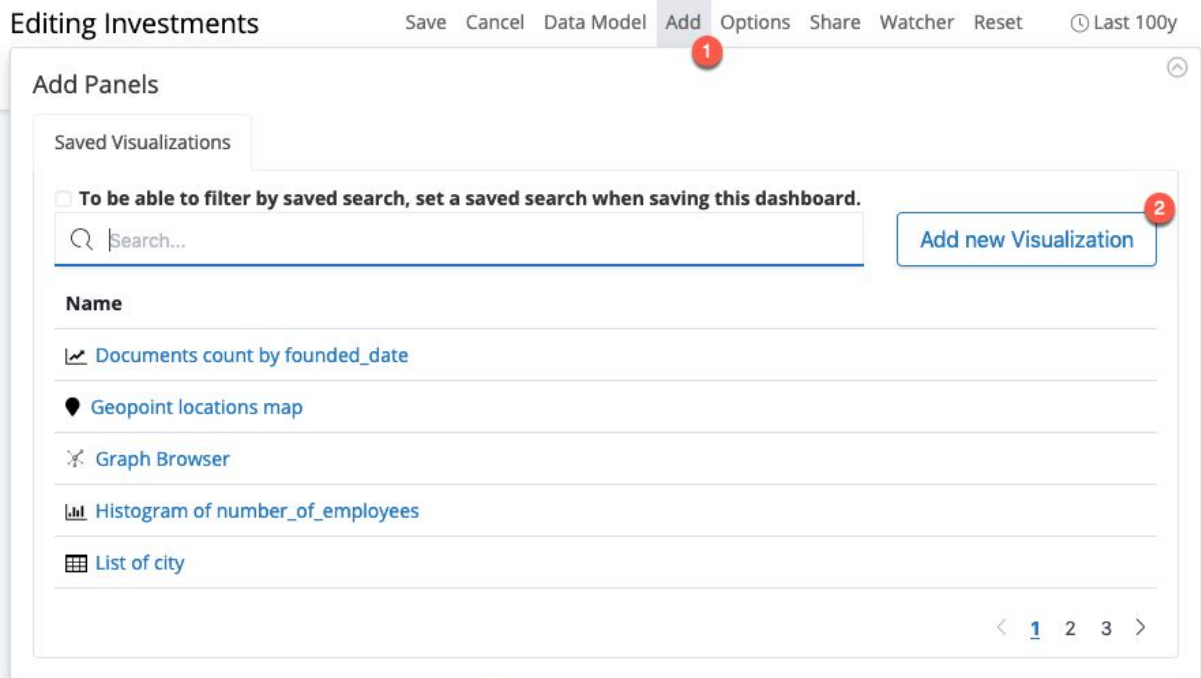
We will add visualizations in the following order:

- Vertical Bar chart
- Record Table

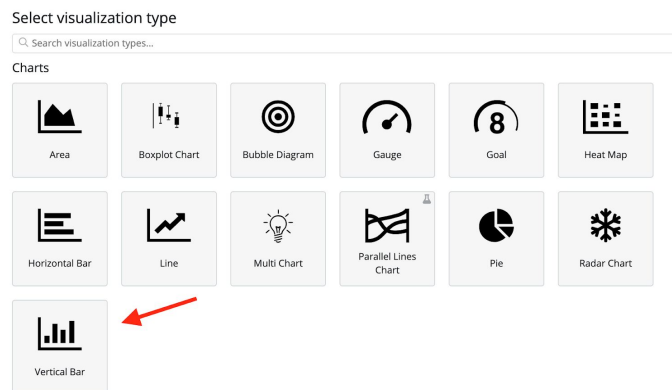
- Analytic Table

Vertical Bar Chart

1. Click **Add** ① from the menu bar at the top.
2. Click **Add new Visualization** ②.

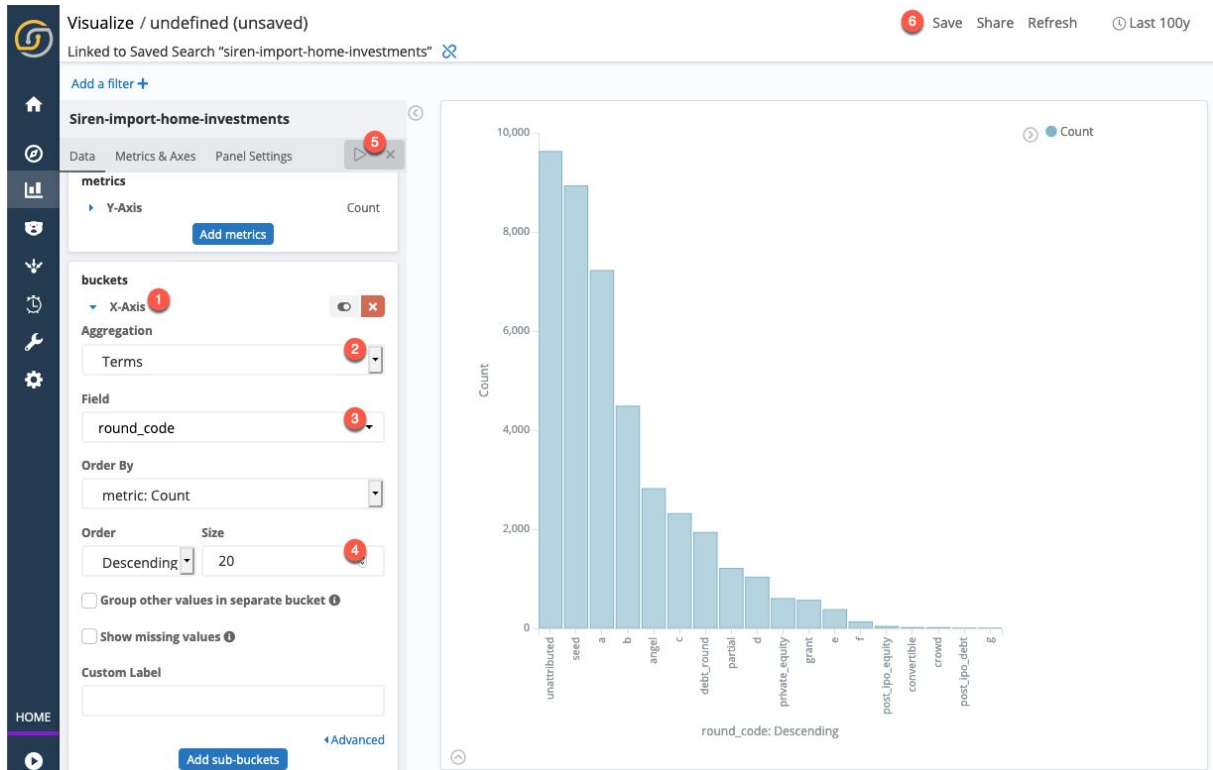


3. Click **Vertical Bar** in the **Charts** category.



4. On the Pick a Saved Search screen, select the **siren-import-home-investments** saved search, since we're building a dashboard for the investments index pattern search.
5. In the **buckets** table window, select **X-axis** ①,
6. In the **Aggregation** field, select **Terms** ②.

7. In the **Field** field, select **round_code** ③
8. In the **Order Size** field, specify a value of **20** ④.
9. Click **Apply changes** ⑤ to see a preview of your chart.
10. Finally, click **Save** ⑥ to save the Visualization.



11. On the **Save** panel, name the visualization “Investment Number by Round Code”, and click **Save and Add to Dashboard**.
12. You are redirected to the **Investments** dashboard in edit mode, where you can resize or reorder the visualizations. Click **Save** to save the dashboard.

Record Table

1. From the dashboard, click **Edit** → **Add** → **Add new Visualization**, as before.
2. Use the search bar or scroll to the Showing individual records category to find **Record Table**.
3. On the **Pick a Saved Search** screen, select **siren-import-home-investments**.
4. Accept the default values in the options fields and click **Apply changes**.
5. Click **Save**, name the visualization “Record Table”, and click **Save and Add to Dashboard**.
6. When you return to the dashboard screen in Edit mode, click **Save**.

TIP: To remove columns that you do not need, click the ‘x’ icon beside the column name. To move columns, use the << >> buttons.

Analytic Table

The Analytic Table visualization is similar in appearance to the Record Table. The difference is that you can view individual records with a Record Table, while you can make summaries with an Analytic Table.

1. From the dashboard, click **Edit** → **Add** → **Add new Visualization**, as before.
2. Use the search bar or scroll to the Textual category to find **Analytic Table**.
3. On the **Pick a Saved Search** screen, select **siren-import-home-investments**.
4. In the **buckets** table window, select **Split Rows**.
5. In the **Aggregation** field, select **Terms**.
6. In the **Field** field, select **raised_currency_code**.
7. In the **Order By** field, select **metric: Count**.
8. In the **Order Size** field, specify the value as **10**.
9. Click **Apply changes**.
10. Click **Save**, name the visualization “Analytic Table”, and click **Save and Add to Dashboard**.
11. When you return to the dashboard screen in Edit mode, click **Save**.

Setting the Dashboard Data Model

(Average completion time: 3 mins)

Before going any further, we need to specify a type of data model for the dashboards. Siren Investigate provides three dashboard options:

- **Dashboard does not represent a search:** Used when the dashboard’s main function is to provide summary information and it does not need to be linked to a search. These dashboards do not display a number (count) or an icon.
- **Dashboard is about a search:** Used when the dashboard represents a single specific search (for example, companies or articles), enabling the following:
 - The dashboard displays the count and an icon associated with the search.
 - Components like Relational Navigator work and can be used on the dashboard.

- The content of the dashboard can be added to a Graph Browser for link analysis.
- **Dashboard 360 with filter strategy:** Uses a *dashboard-specific data model*, which enables a single dashboard to contain visualizations that are based on different searches, and to perform coherent filtering across all of them. This type of dashboard depends on Siren's relational capabilities and the creation of a data model. We will be exploring this option later in the tutorial.

At this point, we have created a dashboard with visualizations that are all linked to the **investments** search. The logical thing to do is to specify that this dashboard is dedicated to the same search, and use the appropriate dashboard data model.

1. On the **Investments** dashboard, click the **Edit** tab, then click the **Data Model** tab
 - ①. You will now see the Dashboard Data Model options:



2. Select the **Dashboard is about a search** option ②, and then select **siren-import-home-investments** from the dropdown list.
3. Click on the **Save** tab ③, and then click the **Save** button.

Importing the investors table

(Average completion time: 5 mins)

But don't create a dashboard yet!

The next file you need for this tutorial is **investors.csv**.

In the **Data Import** app, use the same procedure as previously with the following settings:

1. On the **Excel/CSV Import** screen, set the **Index scope** to **Local** and click **Import**.
2. On the **Create an Index Pattern Search** screen:
 - a. In the **Time Filter field name** field, select **I don't want to use the Time Filter**, as this is not a time-based index.
 - b. Choose an appropriate icon for investors and click **Create**.
3. On the **investors (Index pattern search)** screen, set the **Label when visualized in the graph browser** to be a **Document Field** and, in the second dropdown menu that appears, select the field **label**.
4. Click **Save** in the top-right corner of the screen.

The **Searches** section of the data model should now contain ***companies**, ***investments**, and ***investors**.

IMPORTANT: Do not create a dashboard yet for investors. We will do that later, after we have defined the associative data model.

The screenshot shows the 'Create index pattern search' configuration page in the Siren Platform. The page title is 'siren-import-home-investors (Index pattern search)'. The left sidebar shows a list of searches: 'siren-import-home-com...', 'siren-import-home-inve...', 'siren-import-home-inve...', and 'siren-import-home-inve...'. The main content area has tabs for 'Info', 'Fields (36)', 'Data (14959)', 'Relations', 'Scripted fields (0)', 'Options', 'Revisions', and 'Data model graph'. The 'Info' tab is active, showing a message: 'No dashboard exists for this Index Pattern Search. Did you know? Siren can create one for you, click on the data tab to get started.' Below this, the 'Name of the search' is 'siren-import-home-investors'. The 'Icon' is 'far fa-user'. The 'Color' is '#663db8'. The 'Label when visualized in the graph browser' is set to 'Document Field' with a dropdown menu showing 'label' and a 'preview' link. The 'Short description' field is empty. The 'Index pattern used by this search' is 'siren-import-home-investors'. The 'Time Filter field name' is set to 'I don't want to use the Time Filter' with a 'refresh fields' link.

Introducing relations in the Data Model

The true power of Siren is enabled by the *associative* part of the data model; defining how your tables are interconnected via relations:

In Siren Platform:

- a *relation* is a labeled link between entity types.
- An *index pattern search* can be seen as an entity type, which means that two index pattern searches can be connected by a relation in the Siren data model.

In the example below from our demo, the Investments table is connected to the Companies table.



We will soon see that these relations are created in the UI. For example, this is how the above relation was defined in the “relations” tab of the Data Model. Note that the definition of a relation typically requires one to put both active and passive verbal forms (for example, Companies *secured* investments versus investments that were *secured by* companies):

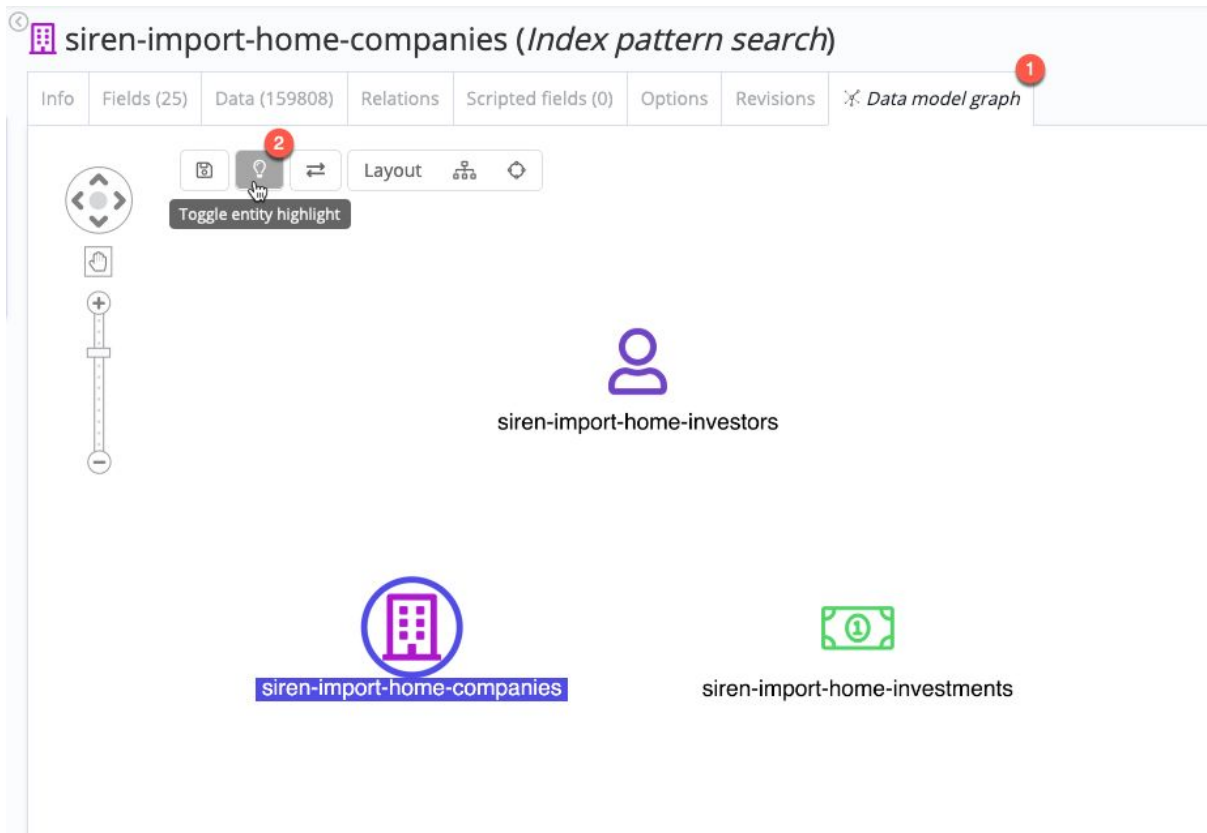
Field	Companies	→	secured	←	Investments	
	id		secured by		companies	

The associative data model is key to exploring the relational data analysis capabilities in the Siren Platform.

It enhances both the navigation across related dashboards and the link analysis feature or visual graph browser.

Creating relations between entities

In the **Management** app, go to the **Data Model** tab. Select the **Data Model Graph** tab ①, which shows the siren-import-home-companies search.



There are three icons, one for each search. To see all of the icons highlighted clearly, click the **Toggle entity highlight** - the lightbulb icon ②. You can see that the indexes are not connected to each other.

We're now going to define the relationship between the indexes and create a relational data model. This can be done manually or by using the auto-discovery wizard.

Creating relations manually

(Average completion time: 5 mins)

In this step, we will work with the **investments** search to specify which fields of other searches relate to it.

Specifically, we will connect the **investments.companies** field with the **companies.id** field, the primary key of companies. We will base this on the [relational model of the data](#).

To create this relation, go to the **siren-import-home-investments** index pattern search, click the **Relations** tab, and complete the following steps:

The screenshot shows the 'Relations' tab for the 'siren-import-home-investments' index pattern search. The interface is divided into three main sections: Source Entity, Labels, and Target Entity.

- Source Entity:** The 'Field' dropdown is set to 'companies' (step 2). Below it, 'investors' is also visible.
- Labels:** The top 'active' field is 'secured by' (step 3) and the bottom 'passive' field is 'secured' (step 4).
- Target Entity:** The 'Field' dropdown is set to 'id' (step 6). Above it, 'siren-import-home-com...' is selected (step 5).

 A blue button labeled 'Add relation +' (step 1) is at the bottom left. A 'Show samples for this relation' button (step 7) is on the right. A 'Relations auto-discovery wizard (BETA)' link is at the bottom right. The top navigation bar includes tabs for Info, Fields (19), Data (41317), Relations (2), Scripted fields (0), Options, Revisions, and Data model graph.

1. Click **Add relation** ①. The source entity is preselected as **siren-import-home-investments**.
2. In the **Field** field, select **companies** ②.
3. In the **Labels** fields, enter “secured by” in the top ‘active’ field ③ and “secured” in the bottom ‘passive’ field ④.
4. In the **Target Entity** field ⑤, select **siren-import-home-companies** from the list of searches.
5. In the **Field** field ⑥, select **id**.
6. Click the **Show samples for this relation** ⑦. The **Relation Explorer** screen shows you sample records from the left and right side, which are matched on the same key. This is a good way to validate that the relationship makes sense. You can move back and forth through documents by clicking the arrows on the top left of the window.
7. Click **Save**. The **Data model graph** tab shows the new relation.

Now, we need to create a relation to the investor. To do this, return to the **Relations** tab and complete the following steps:

1. Click **Add relation**.
2. In the **Field** field, select **investors** from the list of searches.
3. In the **Labels** fields, enter “made by” in the top field and “made” in the bottom field.

4. In the **Target Entity** field, select **investors** from the list of searches.
5. In the **Field** field, select **id** from the list of fields.
6. Click **Save**.

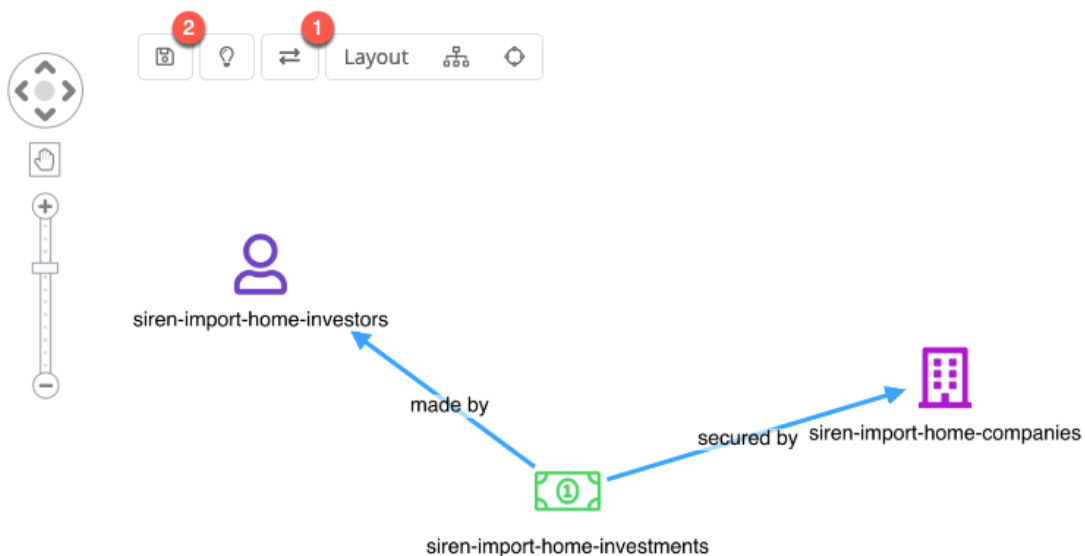
Visualizing the data model as a graph

(Average completion time: 1 min)

To see all the relations in the data model at a glance, click on the **Data Model Graph** tab.

From this screen, you can:

- Click and drag the icons to rearrange them. Change the direction of the relation labels by clicking **Toggle relation direction** ①.
- **Save the graph layout** ②.



Creating a self-relation

(Average completion time: 2 mins)

A self-relation is a relation between two entities of the same type, for example, a person who is *friends with* other people or companies that are *competitors of* other companies. These are defined in the same way.

In our data, the company records have a field called **one_competitor** where you can find the ID of the main competitor of a company. Let's create a self-relation from **companies.id** to **companies.one_competitor**:

1. On the **Data model graph** tab, double-click the **siren-import-home-companies** search icon. This switches to the **siren-import-home-companies** saved search.
2. Click on the **Relations** tab. You'll notice that the relation between companies and investments is already there. Makes sense, right? Now, we'll create a new one - this time **a self-relation**.
3. Click **Add relation**.
4. Specify a relation from **companies.id** to **companies.one_competitor** and enter "competitor" as the label in both directions.
5. Click **Save**.

With the self-relation, the data model graph should now look something like this:



That's it! Your relational data model is done for now. It's time to see it in action.

By the way, did you know that Siren could have built this data model automatically with the relations auto-discovery wizard? Learn more about it in [Appendix 7](#).

Dashboard-to-dashboard associative navigation

The **Relational Navigator** is a special visualization that automatically shows links to connected records in other dashboards. By clicking on these buttons you perform an *associative* navigation, going from one set of records to another set of connected records.

For these buttons to appear in a dashboard, we must add a **Relational Navigator** visualization, the same way we add all visualizations to dashboards.

NOTE: You can create a single **Relational Navigator** visualization and reuse the same visualization across all of the dashboards. Once it is added to a dashboard, it will show buttons that connect with other dashboards that are ‘relationally’ connected.

The auto-generation wizard adds this visualization if - and only if - you have defined the relations between the data. Otherwise, you can add it manually.

Since we have already defined the relations for ‘investors’, we can create a dashboard for it by using the wizard as follows:

1. Go to **Management** → **Data Model**
2. Click on the **siren-import-home-investors** search.
3. Follow the steps to [auto-generate the dashboard](#), specify the title of the dashboard as “Investors” and click **Create**.

You can now see that the **Relational Navigator** visualization, labelled **Related dashboards**, is added to the **Investors** dashboard.

The blue link in this visualization links to the **Investments** dashboard and can be used to find all of the investments that were made by the *investors* on the current dashboard.

Investors (14,959)

Filters
Search... (e.g. status:200 AND extension:PHP)

Related dashboards

[Made \(28723 Investments\)](#)

List of affiliati

affiliation_r
Unaffiliated
Google
Facebook
LinkedIn

- Filtering the records on the dashboard will automatically show a filtered count on the relational navigator link. For example, if you click on the 'person' section of the **Top 30 investortype** pie chart. The record count on the dashboard and the count of investments are reduced as a result.
- Click on **Made (4538 investments)** to go to the **Investments** dashboard. Due to the relational filter that is set, you will see *only investments that are associated with investors of type 'people'*.

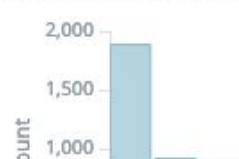
Investments (4,538)

Filters
Search... (e.g. status:200 AND extension:PHP)

coming from (5644) docs on: Investors
Add a filter +

Investment Number by Round Code

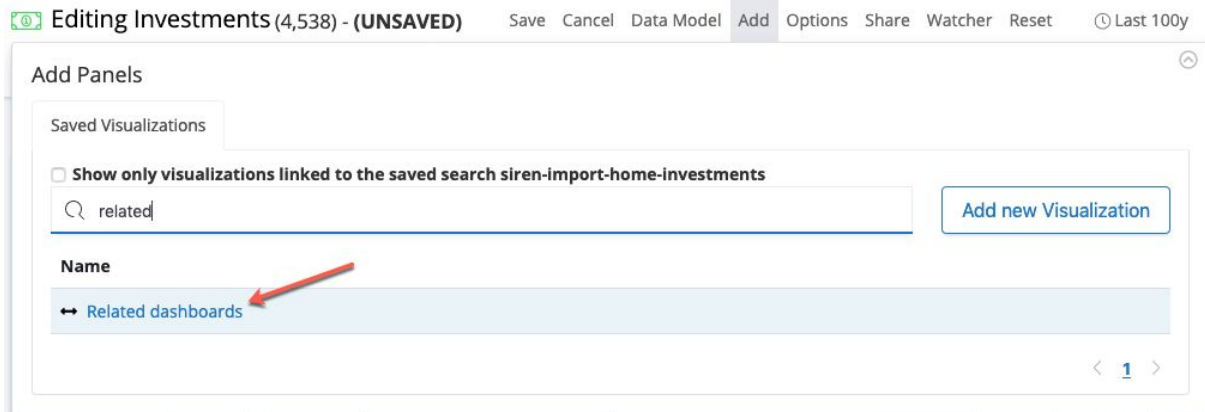
Count



Now, we will add the **Relational Navigator** to the **Investments** dashboard to find more connected records.

You can do this by going to the dashboards that do not have this visualization yet (**Companies** and **Investments**), click **Edit** -> **Add**.

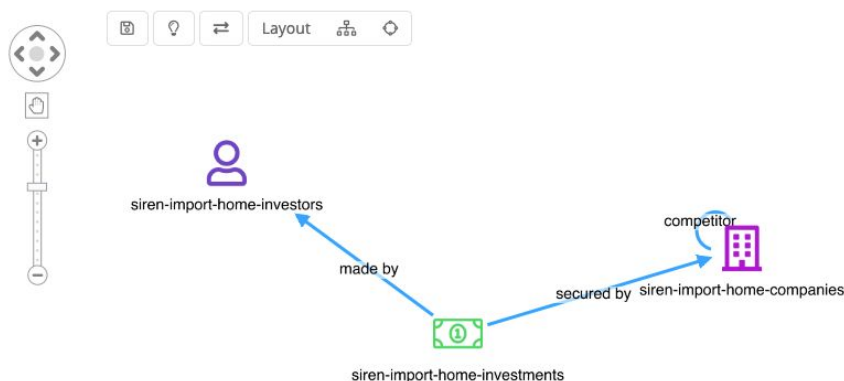
The **Related dashboards** visualization already exists in the system. Therefore, you can simply search for it in the list of existing visualizations, click to add it to the dashboard (a message indicates that it was added to the dashboard successfully), and click **Save**.



TIP: When you are adding panels to dashboards, it is easy to also save filters by mistake. Either remove the filters when you are saving a dashboard or select **Don't overwrite currently saved filter and text query** in the **Save** panel.

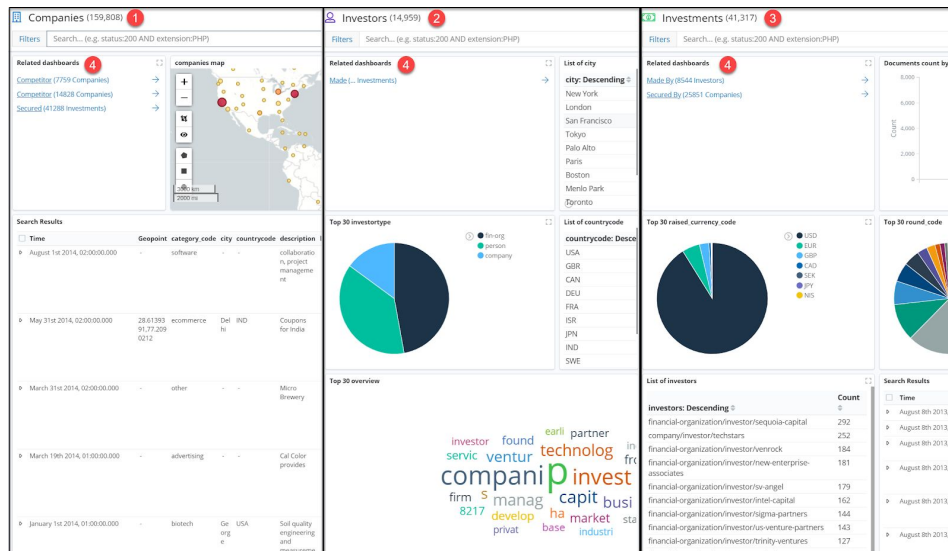
Recap: Where are we now?

If you completed all of the previous steps of this tutorial successfully, your data model should look like this:



You should also have created three dashboards: Companies ①, Investors ② and Investments ③. The number shown after each dashboard name represents the number of documents in the relative Index Pattern Search.

Most importantly, each dashboard must contain a Relational Navigator visualization, called **Related dashboards** ④. You can rearrange the visualizations to bring this dashboard to the top.



If everything looks correct, we can now proceed to explore the data in more detail by completing some exercises on each dashboard and visualizing the results in the **Graph Browser**.

EXERCISE: Investigating investments

In this section, we will answer a few advanced questions by analyzing the data.

NOTE: At any time, you can click the **Reset** button to restore the dashboard to its saved state. Alternatively, you can click the **Reset all** button to restore the state of all of the dashboards.

This exercise contains four parts, in which we will:

1. Answer a basic question by using the **Investors** dashboard.
2. Answer a question that involves a *relational navigation*, from **Investors** to **Investments**.
3. Navigate further to the **Companies** dashboard.
4. Use a link analysis to explore the results together as a single picture.

Part 1: Finding an answer by filtering the Investors dashboard

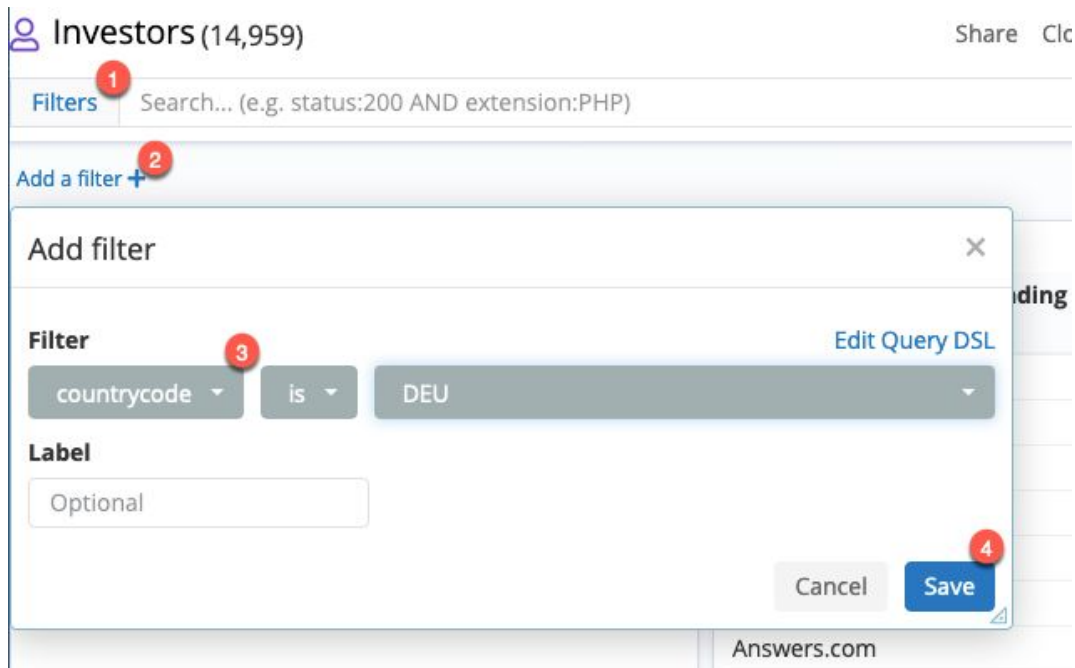
Question: How many investors are from Germany?

Answer:

Basic solution

To find the answer to this question:

1. Go to the **Investors** dashboard.
2. Click **Filters** ① -> **Add a filter** ②.
3. Select **countrycode**, **is** and **DEU** (the iso code for Germany) from the drop-down menus ③.



4. Click **Save** ④. The answer appears as the number in parentheses alongside the dashboard name.

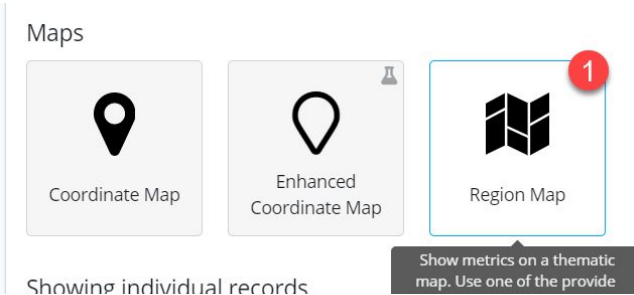
Advanced solution (optional)

You can also answer this question by adding a **Region Map** visualization.

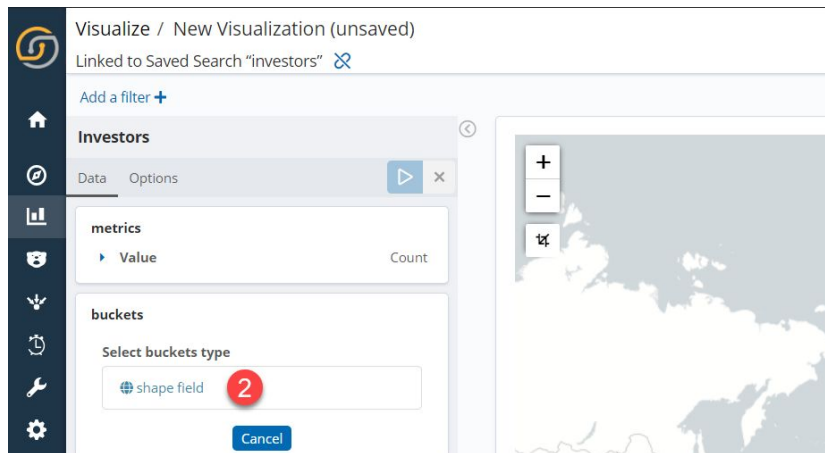
A **Region Map** is a clickable map that allows users to create filters in the dashboard in an interactive fashion. Unlike the **Enhanced Coordinate Map** visualization that works using

geopoints (latitudes and longitudes stored in the records), the Region Map works by matching field values, such as **countrycode:deu**, with shapes and their names.

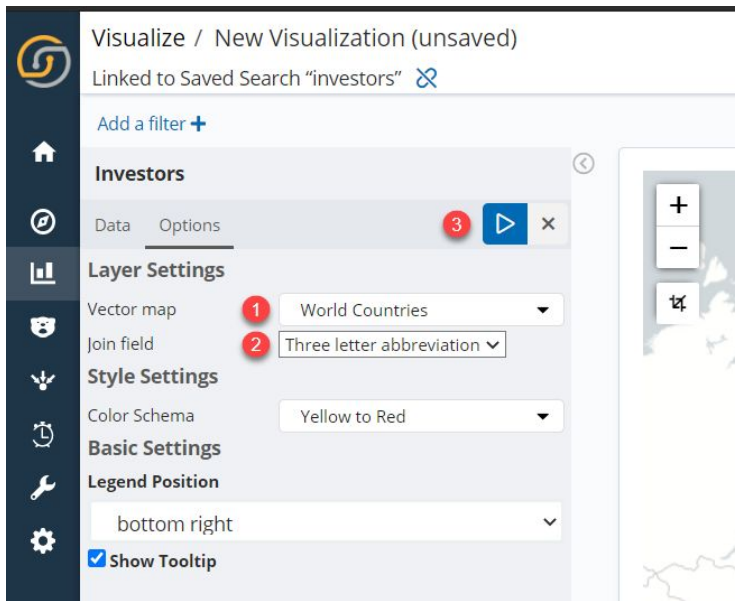
1. As previously shown in the [Creating visualizations](#) section, click **Edit** -> **Add** -> **Add new visualization**. Under the map category, choose **Region Map** ①.



2. In the **Pick a saved search** window, select the **siren-import-home-investors** search.
3. In the **buckets** table window, select **shape field** ②.



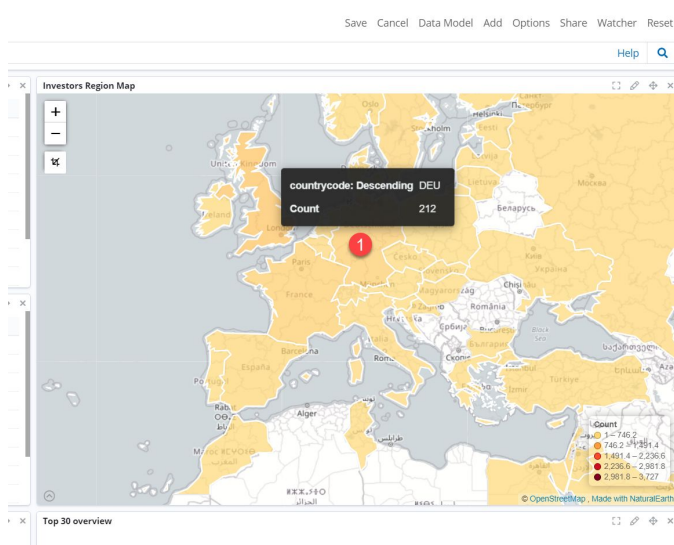
4. In the **Field** field, select **countrycode**.
5. We want to see data from more than five countries, so in the **Size** field, increase the value to **35**.
6. Click the **Options** tab to select the proper layer settings to match the region type of our data.
7. In the **Vector map** dropdown menu, select **World Countries** ① and in the **Join** field, select **Three letter abbreviation** ②.
8. Finally, click **Apply changes** ③.



9. Click **Save**, name the visualization “Investors Region Map”, and click **Save and Add to Dashboard**.

10. Move the Region Map into the position you want and save the **Investors** dashboard.

HINT: Hover your cursor over Germany. The **count** value appears in the tooltip.



To create a filter, click on Germany (DEU). The dashboard count will change accordingly and give you your answer.

Q: How many investors are from Germany?

A: 212

Part 2: Finding an answer by using relational navigation

Question: How much money did German investors invest between 2010 and 2012?

Answer:

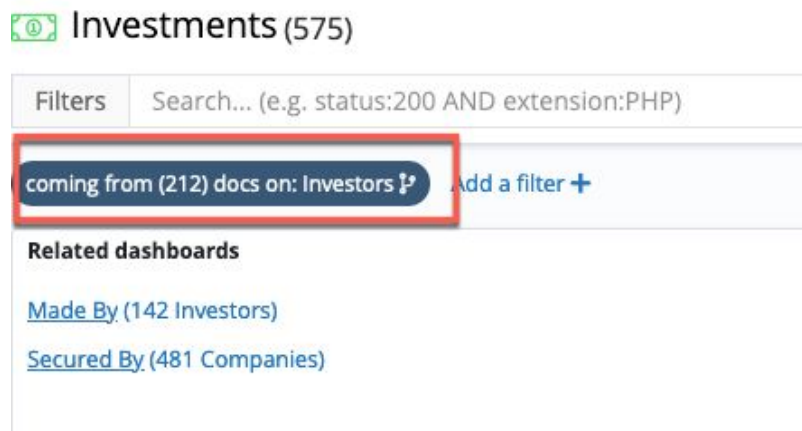
Solution

Now that we have only German investors in the dashboard, we can use the **Relational Navigator** to go to the **Investments** dashboard.

The button already shows you the number of target records: Out of 41k unfiltered investments, only 575 were made by German investors.

Let's go ahead and click on **Made (575 Investments)**.

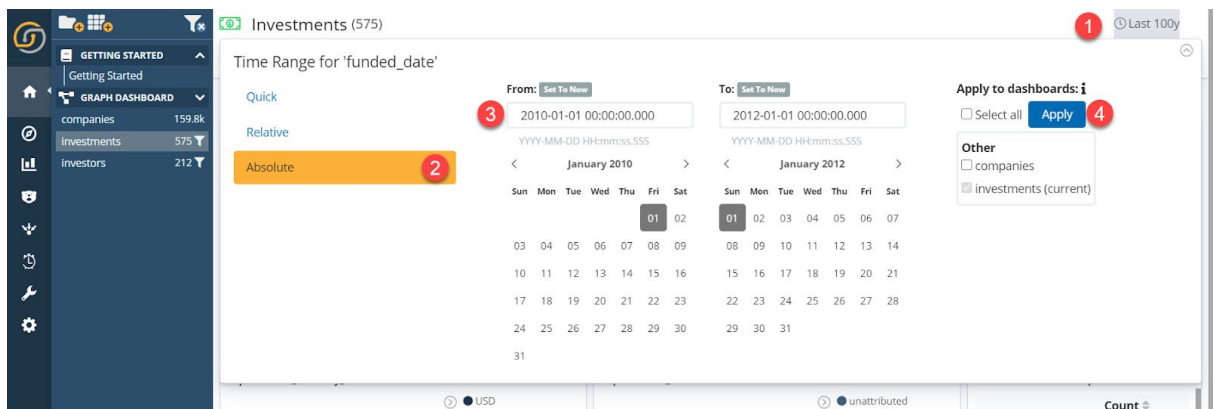
On the **Investments** dashboard, a relational filter is automatically created:



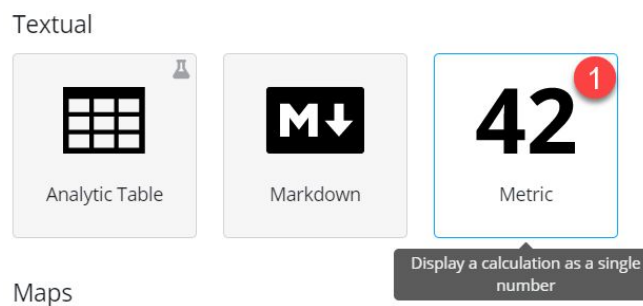
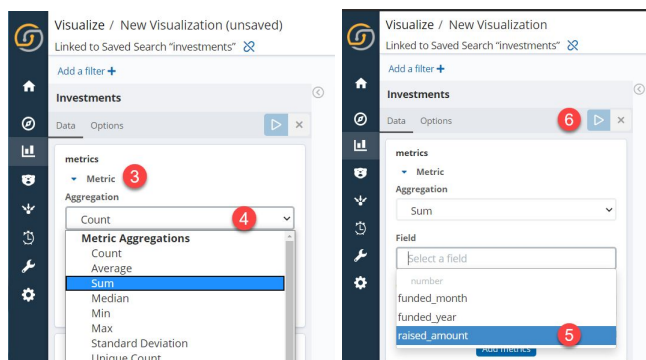
However, we cannot find the answer to the question yet, because we are missing a filter for the two-year time window and a metric to compute the total funding amount.

Let's fix that!

1. To change the time window, click **Last 100y** ①.
2. In the **Time Range** panel, select **Absolute** ② and specify the desired time span ③ (2010 - 2012).

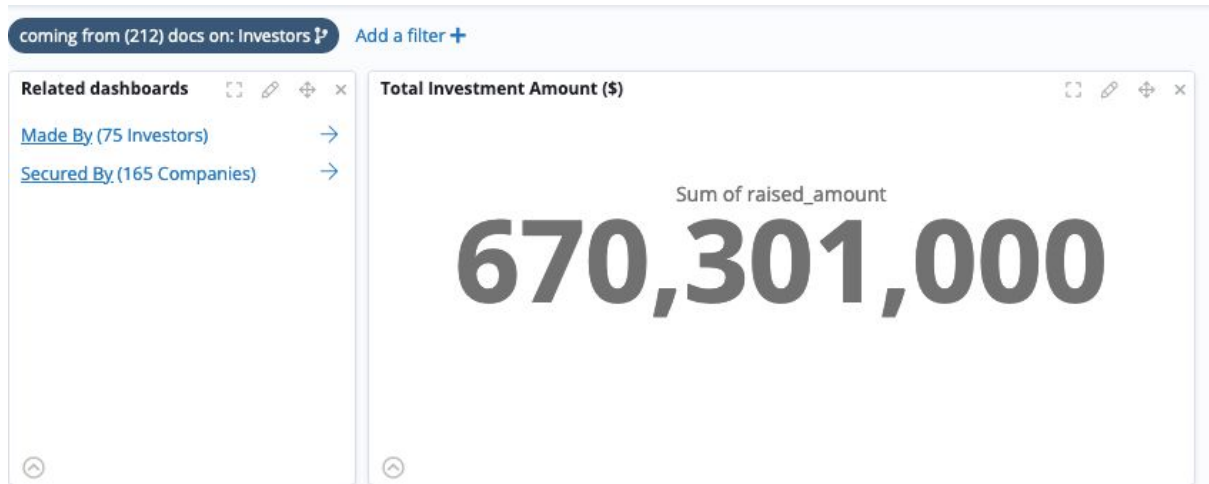
3. Click **Apply** ④.

Perfect! We only need one last step to answer this question, which is, adding a metric to compute the total amount of the 176 investments that are left.

4. Click **Edit** -> **Add** -> **Add new visualization**. Under the Textual category, choose **Metric** ①.5. From the **Pick a saved search** menu, select **siren-import-home-investments**.6. In the visualization editor, expand the **Metric** table window ③.7. In the **Aggregation** field, select **Sum** ④.8. In the **Field** field, select **raised_amount** ⑤.9. Click **Apply changes** ⑥.10. Save the visualization as "Total Investment Amount (\$)" and add it to the **Investments** dashboard.

11. Move the new visualization into position and save the dashboard. This final step will provide us with the answer to the exercise question and you'll permanently have access to this useful metric.

Q: How much money did German stockholders invest between 2010 and 2012?



A: German stockholders invested ~670M \$ between 2010 and 2012.

Part 3: Digging deeper with relational navigation

Question: How much of these investments were secured by American companies?

Answer:

Solution

To solve this question, complete the following steps:

1. Click on the **Companies** dashboard in the sidebar (do not navigate there by using the relational navigator). This will show you all of the companies (approx. 159k).
2. Create a filter for country code **USA** by clicking on **Filter for value** in the **List of countrycode** visualization.

List of countrycode

countrycode: Descending	Count
USA	46,033
GBR	6,175
IND	3,129
CAN	3,049
DEU	1,601
FRA	1,398
AUS	1,138
ESP	918
ISR	852
NLD	738

Export: [Raw](#) [Formatted](#)

1 2 3 4 5 ...10 »

Now, we have selected only American companies.

The **Relational Navigator** is showing us that **34** of the records in the **Investments dashboard** are matching. Notice that, at this point, the Investments dashboard is already filtered from our previous exercise (it contains the filter that displays only the investments by German investors).

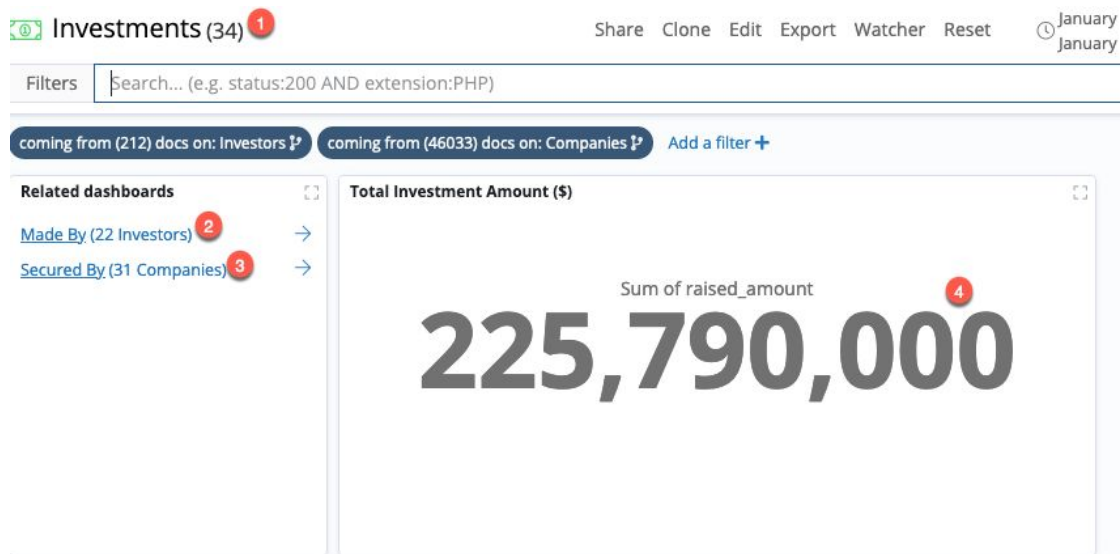
So, let's click on **Secured (34 investments)** to reveal the investments.

countrycode: "USA" Add a filter +

Related dashboards	
Competitor (4965 Companies)	→
Competitor (2554 Companies)	→
Secured (34 Investments)	→

You will be brought to the **Investments** dashboard again but, this time, there are two relational filters present:

- One filter indicates that we are looking at investments “from German Investors” (212);
- The other filter indicates that they were “secured by American companies” (46033).



As you can see in the above screenshot, these 34 investments ① - made by 22 German investors ② and secured by 31 American companies ③ - come to a total of ~\$225M ④.

Q: What amount of these investments were secured by American companies?

A: ~225M \$

Part 4: Link analysis in the Graph Browser

Exploring the results with the Graph Browser

Now that we have solved all of the questions, it is time to explore the results by using the link analysis feature in the **Graph Browser** visualization.

Before doing so, let's apply the relational filter on the other dashboard to focus our research around German Investments, Investors and American Companies.

1. Click **Secured By (31 Companies)** and take a moment to see who these companies are in the relative dashboard.

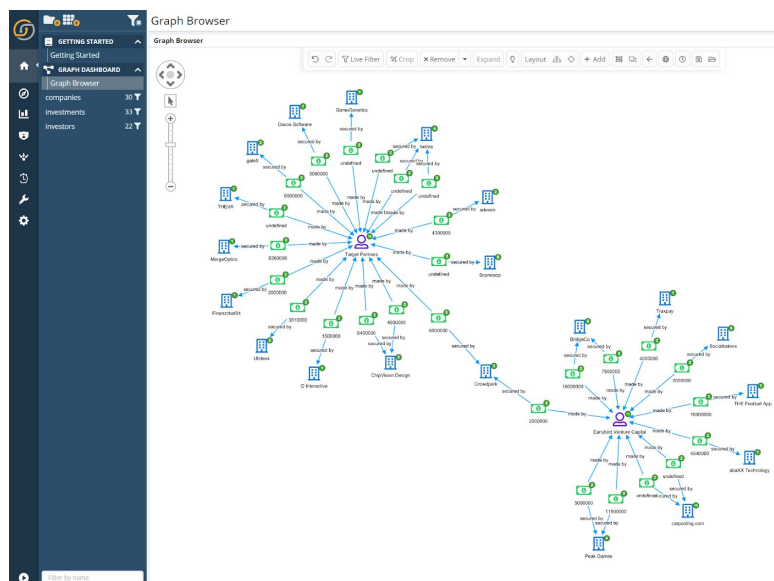
- When you are ready, go back to the **Investments** dashboard and, this time, click **Made By (22 Investors)**.

If everything was done correctly, your filtered dashboards should now display:

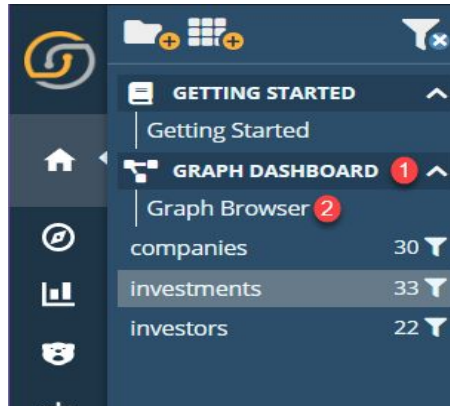
- 31 Companies (all of them American).
- 34 Investments (made by German investors and secured by the 30 American companies)
- 22 Investors (German stockholders that invested in american Companies)

Let's explore together how the Companies, Investments and Investors that we found in our small investigation are linked together and if there are some interesting common connections.

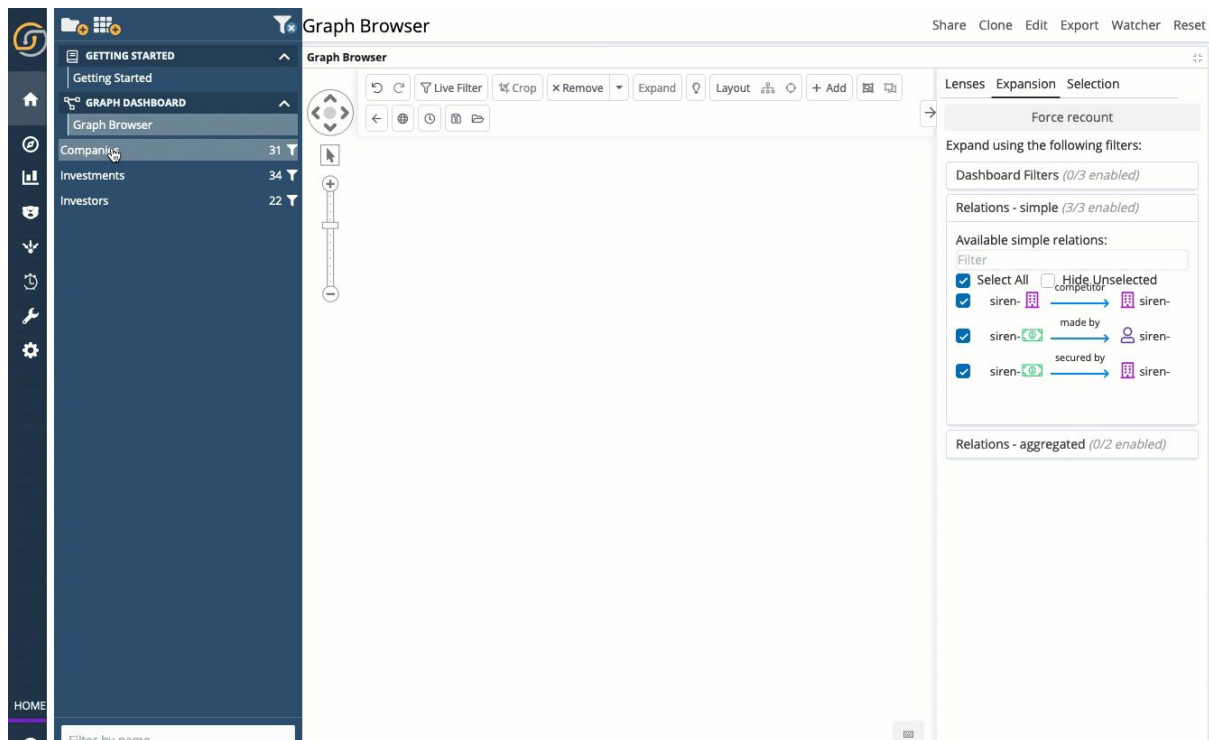
In this part, we'll get to see graphs like these:



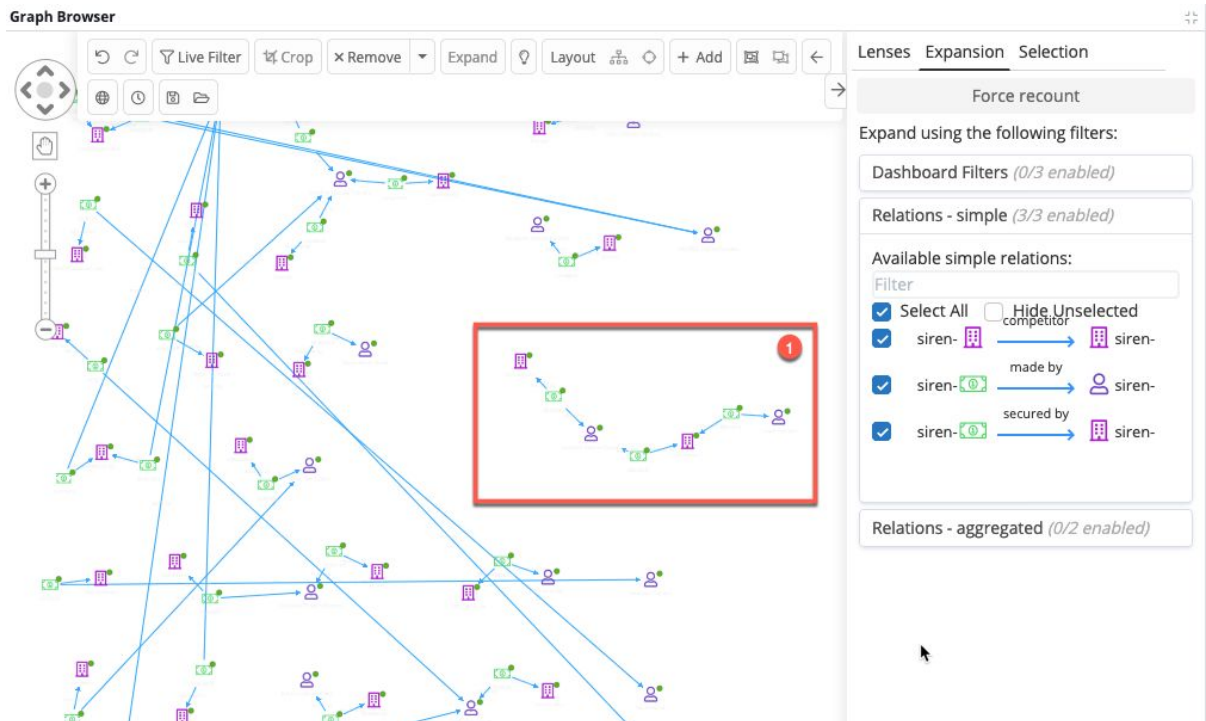
TIP: The Easy Start edition comes with a preconfigured Graph Browser, where you can perform Link Analysis.



1. To start exploring the data, expand the **GRAPH DASHBOARD** group ① and click **Graph Browser** ②.
2. From the Graph Browser visualization, simply drag and drop the **companies**, **investments** and **investors** dashboards into the graph.

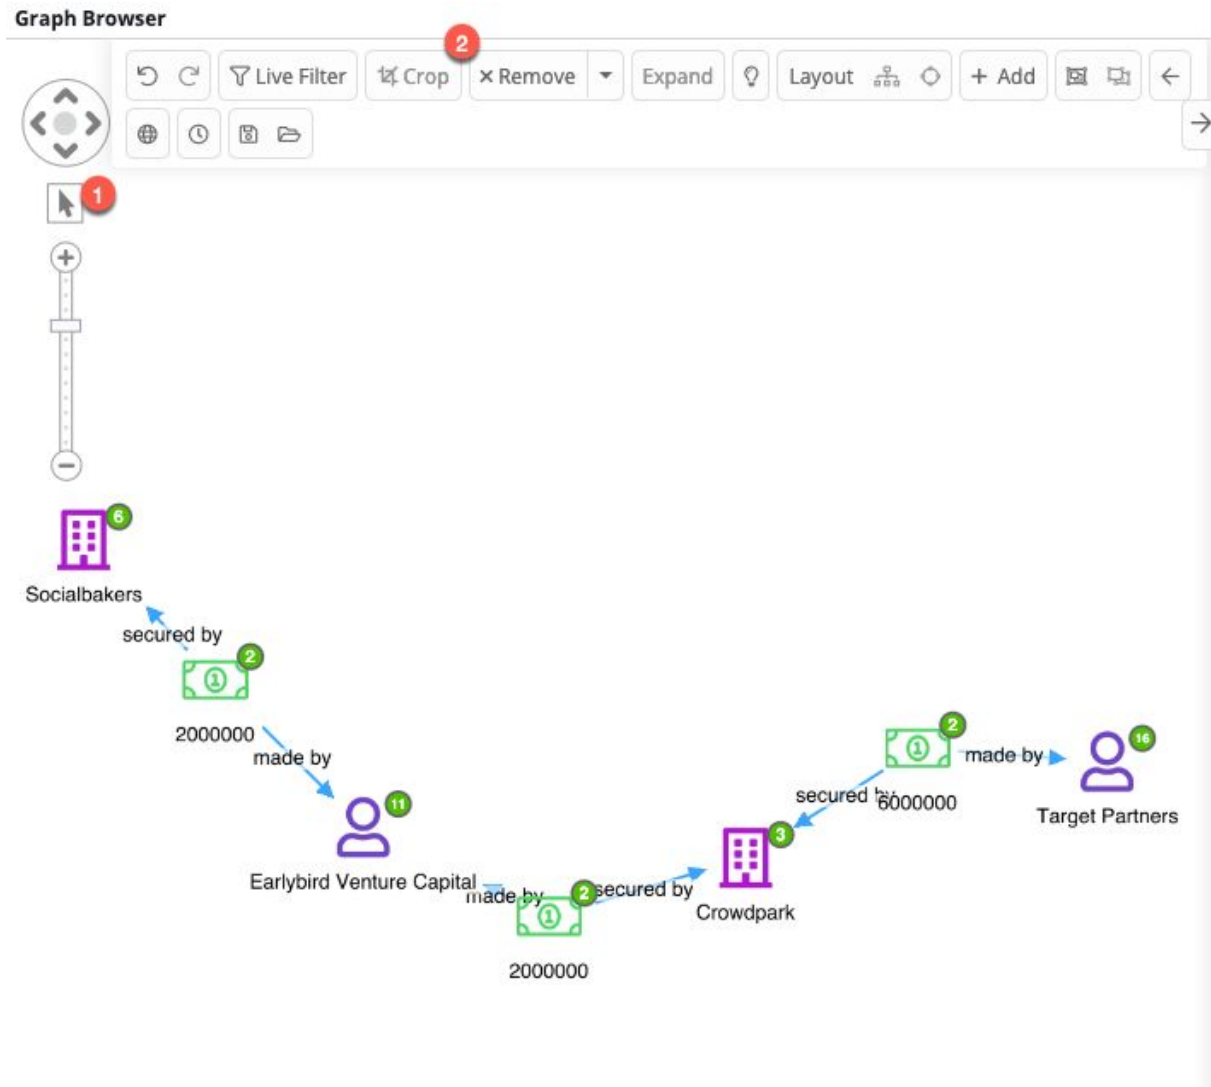


3. Click **Layout** and you should see something like this:



Interestingly, it seems that 2 German stockholders invested in the same company ①. Zoom in on the graph to see this network, which features a company named *Crowdpark*. Let's explore further this small but interesting network.

4. Set the pointer to selection mode ①, drag a rectangle around the network that was identified in the previous screenshot and click **Crop** ②. This will remove everything else from the Graph Browser.



5. Before expanding our network, it is good practice to save this initial state. To do so, click **Save graph** (the disk icon, in the top-right of the menu bar) and save the graph with a suitable name, such as “Shared investment”.

Now that we have saved our initial network, we can start asking more interesting questions.

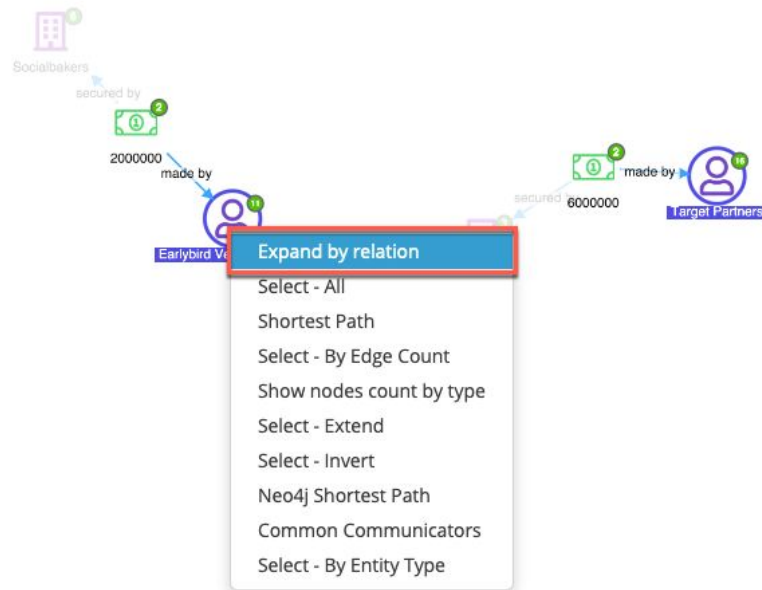
Before we proceed with the tutorial, take a moment to consider the links that might be interesting to explore. Any ideas?

We selected this network because the two German investors share an American company in their investment campaign. It would be interesting to uncover that:

1. They have more common investments.
2. They invested in CrowdPark (the funded company) in the same time period.

Let's start by investigating the first point:

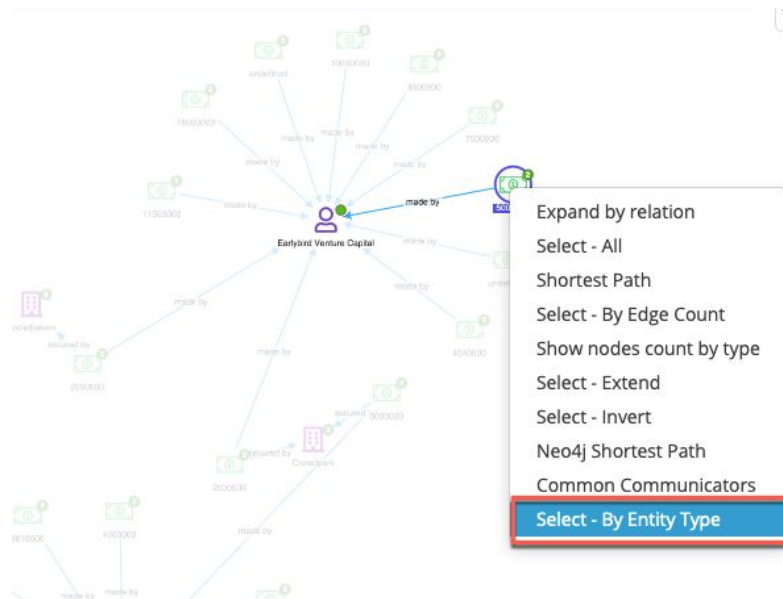
1. While in selection mode, hold the **Ctrl** key on the keyboard and click on the two investor nodes.
2. Once the two nodes are selected, right-click on one of them and, from the contextual menu, select **Expand by relation**.



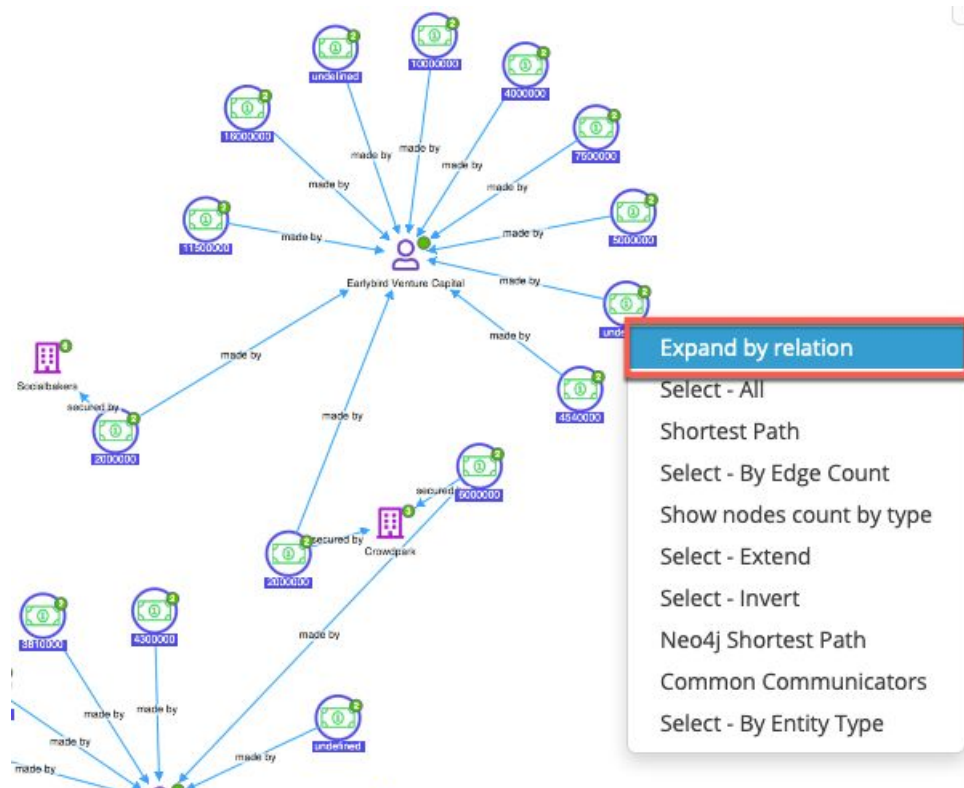
3. Select the **made (27) siren-import-home-investments** checkbox and click **Ok**.

We will now investigate whether one of the new investments that appeared is paid to a new common company.

4. To do so, right-click any one of the investment nodes and click **Select - By Entity Type**.



5. In the pop-up window, select **siren-home-import-investments** and click **Ok**.
6. Now, all of the investment nodes should be selected. Right-click on one of the nodes and select **Expand by relation**.

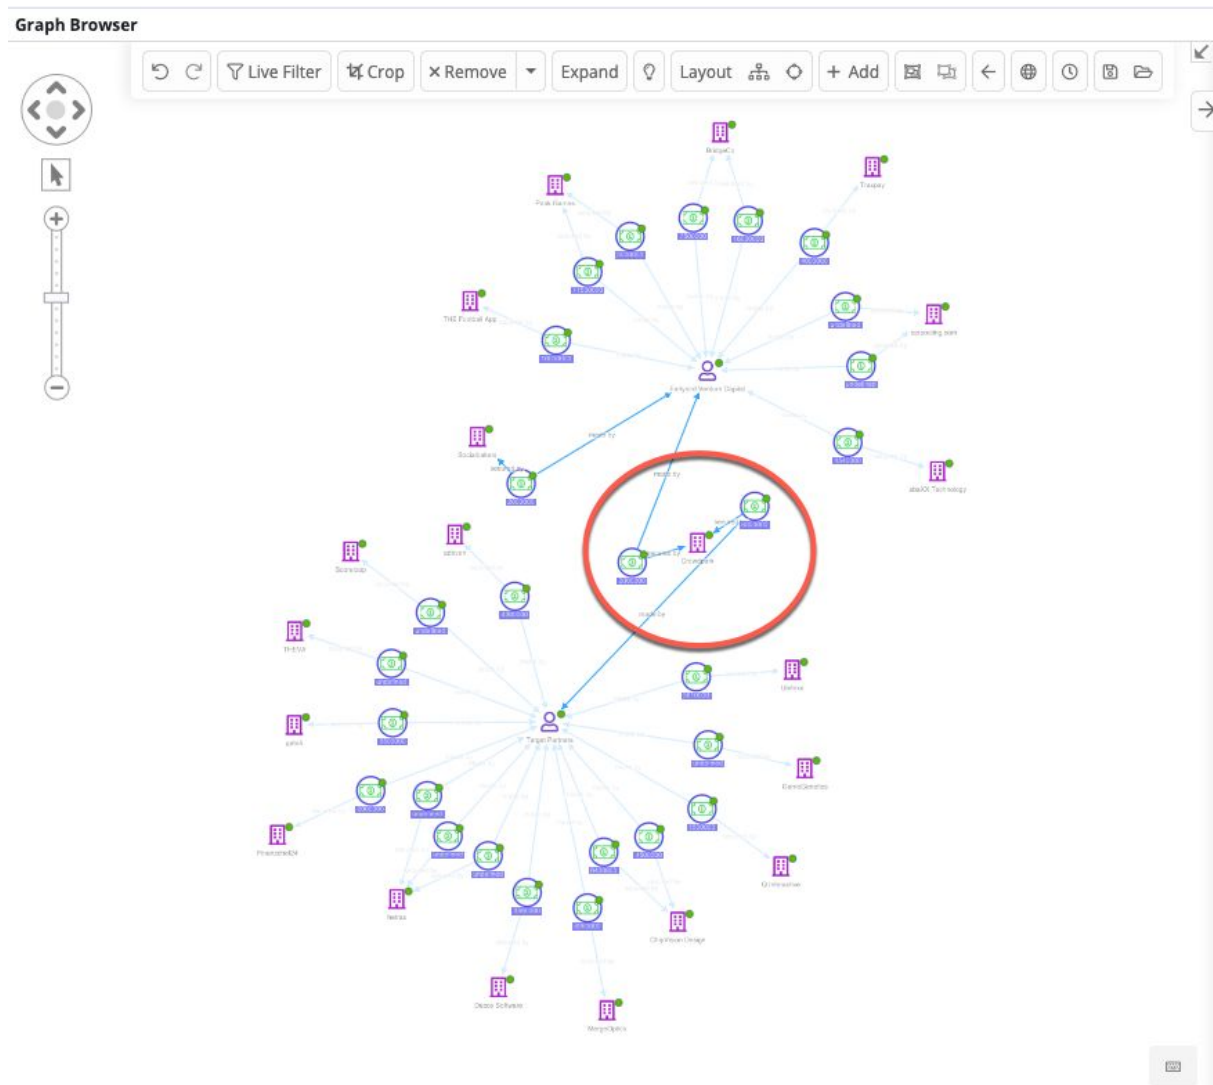


7. In the pop-up window, select **secured by (27) siren-home-import-companies** and click **Ok**.

As we can see in the graph below, the only common connection between the two German investors (from 2010 to 2012) is *CrowdPark*.

Q: Do the two German stockholders share additional companies in their funding campaign?

A: No.

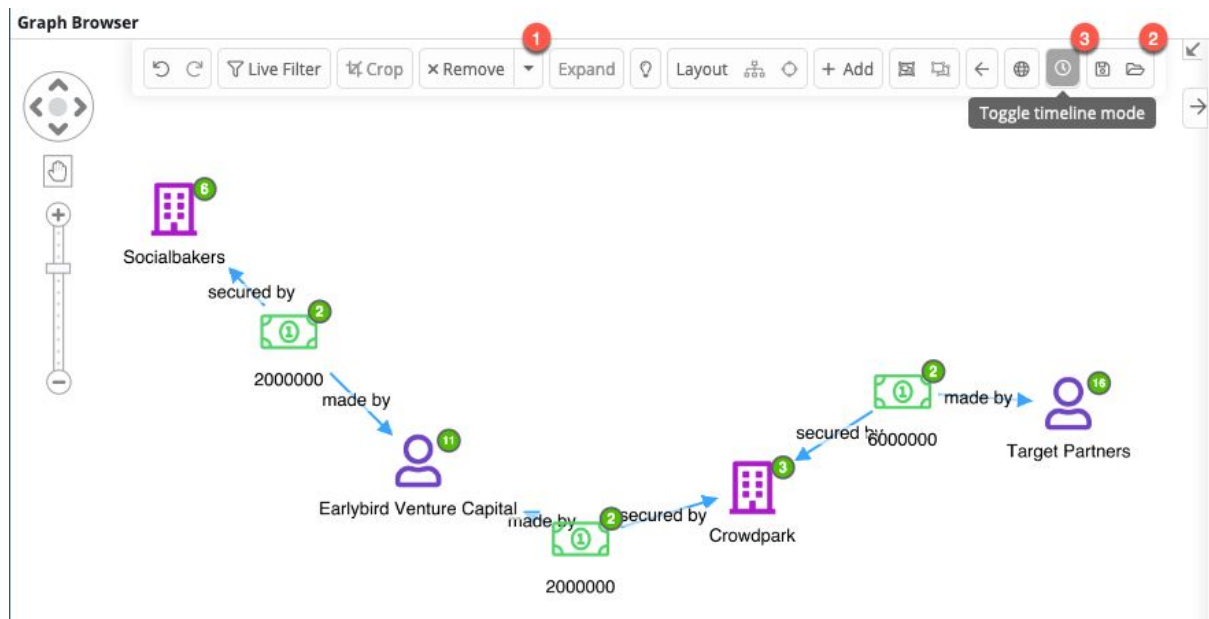


NOTE: If you notice that the labels of the entities are not displaying correctly, ensure that you completed all of the steps to set the labels in your index pattern searches, as described, in steps 7-8 of [Creating your first index pattern search](#).

TIP: Before you proceed further in the tutorial, feel free to expand and explore the entities further and test the different functions that are available in the Graph Browser menu bar.

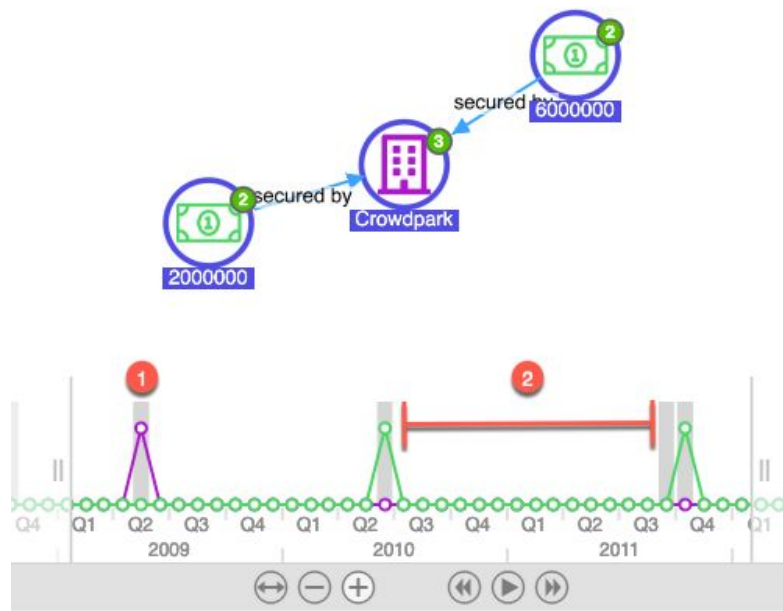
The second interesting question was: Did the two German investors fund CrowdPark in the same time period?

1. Remove the nodes and return the graph to its saved state by clicking **Remove All** ①.
2. Click **Open saved graph** ② and select the saved graph from the **Load graph** window.
3. Activate the timeline mode functionality of the graph by clicking **Toggle timeline mode** ③.



In timeline mode, hold the **Shift** key and select the nodes that we are interested in; the company CrowdPark and the two investments. This will allow us to explore the temporal dynamics of the selected nodes.

As we can see in the screenshot below, CrowdPark was founded in the second quarter of 2009, while the two investments were made approximately one year apart.



Q: Did the two German investors fund CrowdPark in the same time period?

A: No.

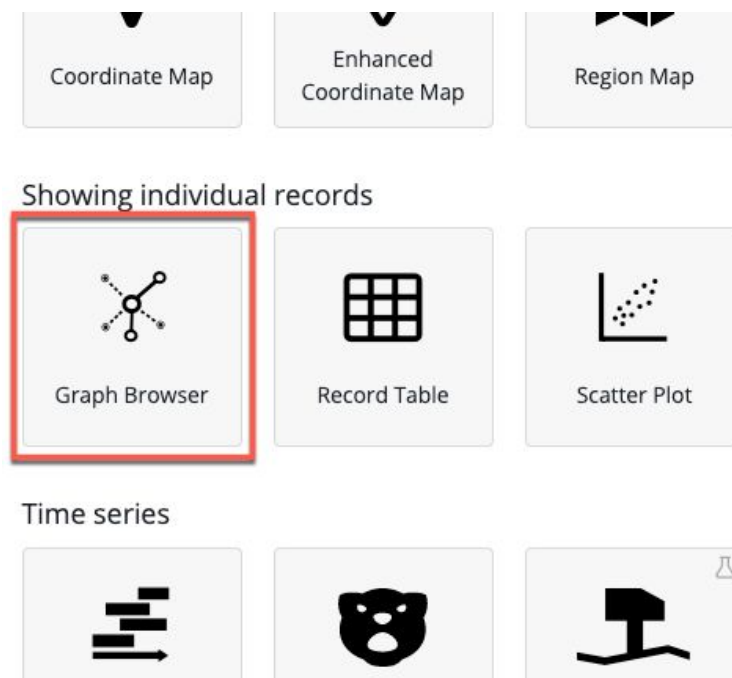
This concludes our exercise! If you want to know more about the graph, we will cover more material in the next chapter.

More about link analysis

Embedded in dashboards

Link analysis is one of the key features and can be performed in the **Graph Browser** visualization.

In the exercise, we used a pre-configured dashboard with a **Graph Browser** component, however, you can add it to a dashboard by selecting it from the **Visualizations** menu:



Using the Graph Browser

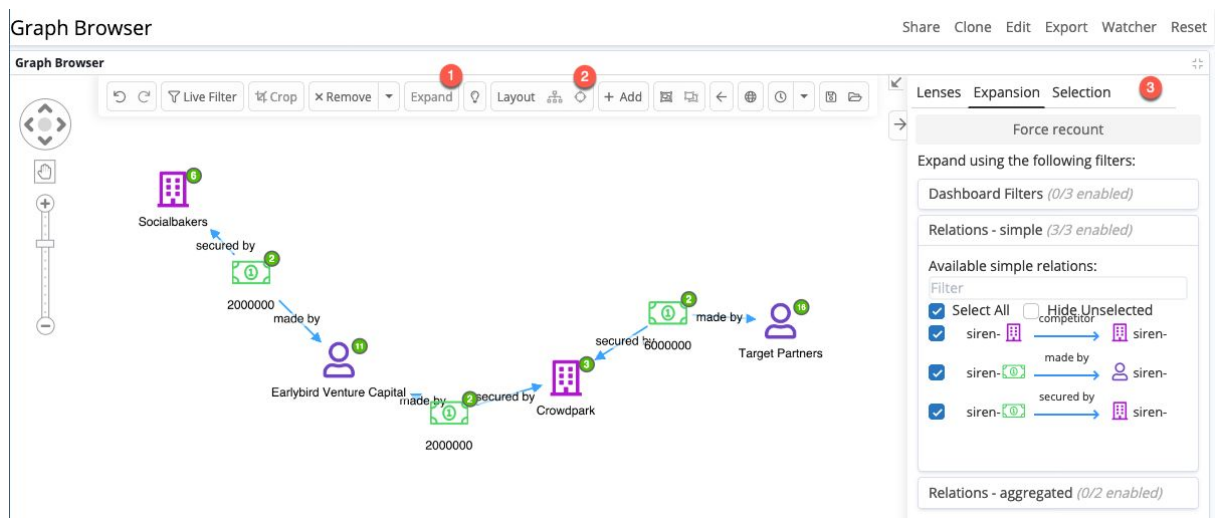
You can add data to the Graph Browser visualization by dragging and dropping filtered or unfiltered data from dashboards.

You can expand the data by double-clicking on any node but also by using the **Expand** button ①, which acts on all of the selected nodes. Clicking the **Layout** button ② reorganizes the nodes.

The Graph Browser also offers some advanced features to help with link analysis.


The menu bar ③ contains the following tabs by default:

- **Lenses** - Create lenses to alter the way the graph looks. For example, create conditional formattings for labels, colors, sizes, or change visibility. For more information about lenses, see the [Siren Platform documentation](#) or follow the steps to create a lens in [Appendix 2](#).
- **Expansion** - Controls the way the graph expands (either when you click the expand button or when you double-click on nodes). Here, you can set filters (by using Dashboard Filters) or choose the data model relation to consider when expanding. For more information, see the [Siren Platform documentation](#).
- **Selection** - See details about the current selection and search within the selection.



The **Cards** tab is not available by default. To configure the **Cards** tab, complete the following steps:

1. In the **Graph Browser**, click **Edit** in the top menu bar.
2. Click the pencil icon ① to open the **Edit visualization** menu.



3. Click the **Scripts** tab and, under the **Cards** section, click **Add card script** and select **Histogram Card - default implementation** from the dropdown menu.
4. Click **Apply changes** and **Save**. You might need to refresh the page for the changes to appear.

For more information about using the **Cards** tab, see the [Siren Platform documentation](#).

Appendix 1: Natural Language Processing (NLP) and more link analysis

The Siren Platform Easy Start version comes with the Siren Platform NLP plugin pre-installed, which adds an NLP *pipeline processor* to enrich documents that are being ingested in Elasticsearch (by a CSV file import or by another method) with entity annotations².

Here, we will use it to enrich the fourth CSV file in our demo dataset, the **articles.csv** file, which contains 640k tech articles that are harvested from the Web.

The pipeline processor will take a text field and produce annotations for Named Entities (Organization, Person, Location), which are then useful to connect to the data that we already have. For example, connecting articles to companies or investors.

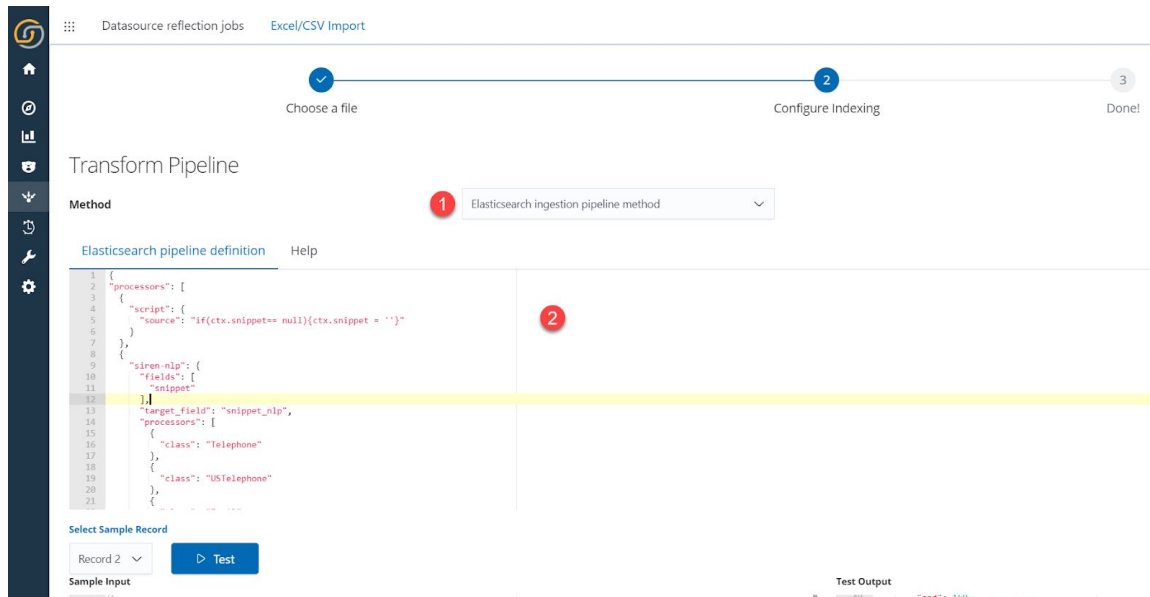
To do this, we will activate a ‘transformation pipeline’ with the NLP processor, while importing data into the **articles** index.

Before you begin

If you are using a slow machine, you can go to **Management** → **Advanced Settings** and edit the value for **ingest-xlsx:bulk_package_size** field to 20.

1. Import the **articles.csv** file that is provided with the [sample data](#) as usual but, this time, in the **Transform Pipeline Method** dropdown menu, select **Elasticsearch ingestion pipeline method** ①.

² It is always possible in Siren Platform to use any other NLP engine as part of ETL, but Siren NLP is made available via the Elasticsearch pipeline processor, which makes it easy to activate and requires no external server or process.



2. Copy and paste the below JSON processor in the **Elasticsearch pipeline definition**

②.

```
{
  "processors": [
    {
      "siren-nlp": {
        "fields": [
          "snippet"
        ],
        "processors": [
          {
            "class": "Telephone"
          },
          {
            "class": "USTelephone"
          },
          {
            "class": "Email"
          },
          {
            "class": "IPv4"
          },
          {
            "class": "IPv6"
          },
          {
            "class": "MacAddress"
          },
          {
            "class": "Url",
            "settings": {
              "lenient": "true"
            }
          },
          {
            "class": "SortCode"
          },
          {
            "class": "HashTag"
          },
          {
            "class": "NER",
            "settings": {
              "nerType": "Organization"
            }
          },
          {
            "class": "NER",
            "settings": {
              "nerType": "Person"
            }
          }
        ]
      }
    }
  ]
}
```

```

    }
  },
  {
    "class": "NER",
    "settings": {
      "nerType": "Location"
    }
  }
]
}
}
}

```

The **nlp-pipeline** is now configured to perform the following tasks:

- Extract annotations by using various processors from the **snippet** field of the **articles** index.
 - Write the annotations to the field **siren.nlp**.
 - Extract the following items using regex or simple rules: Telephone, USTelephone, Email, IPv4, IPv6, MacAddress, Url, SortCode, HashTag.
 - Extract Organization, Location and Person Named Entities by using OpenNLP [NER models](#).
3. Under **Mapping**, select **Date** as the type for the **pdate** field and **Text (Allows Word Cloud)** as the type for the **snippet** field.
 4. Click the **Import** button. The import might take about 30 minutes.
 5. Click **Create an Index Pattern Search** and use **pdate** as the **Time Filter field**.

Connecting extracted entities in the data model

Now, let's define the relationship between the **articles** index and the entities and the rest of the indexes in our data model. To create relations manually, complete the following steps:

1. Go to the **Relations** tab and click **Add relations**.
2. Under Source Entity **articles**, select **companies** as the Field.
3. For Target Entity, select **companies** and select **id** as the Field.
4. Add the labels, **mentions** and **mentioned in**, as shown in the screenshot below.
5. Click **Save**.

Example usage of NLP Data in Siren Investigate

Data annotated by using the NLP plugin can be very useful for finding interesting relationships between fields in an index and the annotated text. In this example, we will try to find the mentions of companies as a part of the articles' snippet.

On the **Data Model** page, click on the **Data** tab. Under the fields, you can see that a new field **siren.nlp** has been created by the NLP plugin during the import. This field contains annotations for fields processed by Siren NLP (in this case the **snippet** field), such as organization, person and location.

From the Data model page for the articles index, complete the following steps:

1. Create a relation between Articles **siren.nlp.matches.entity/organization** and Companies **label**.

articles (*Index pattern search*)

Info Fields (195) Data (646902) Relations (2) Scripted fields (0) Options Revisions [Data model graph](#) ☆

Source Entity	Labels	Target Entity
Field: articles companies.keyword	→ mentions ← mentioned in	Field: companies id
Field: articles siren.nlp.matches.entity/organiz...	→ mentions ← mentioned in	Field: companies label

[Add relation +](#) [Relations auto-discovery wizard \(BETA\) >](#)

2. Given that we don't have an index for "Persons", [create an Entity Identifier \(EID\)](#) and add a relation from it to **siren.nlp.matches.entity/person**.

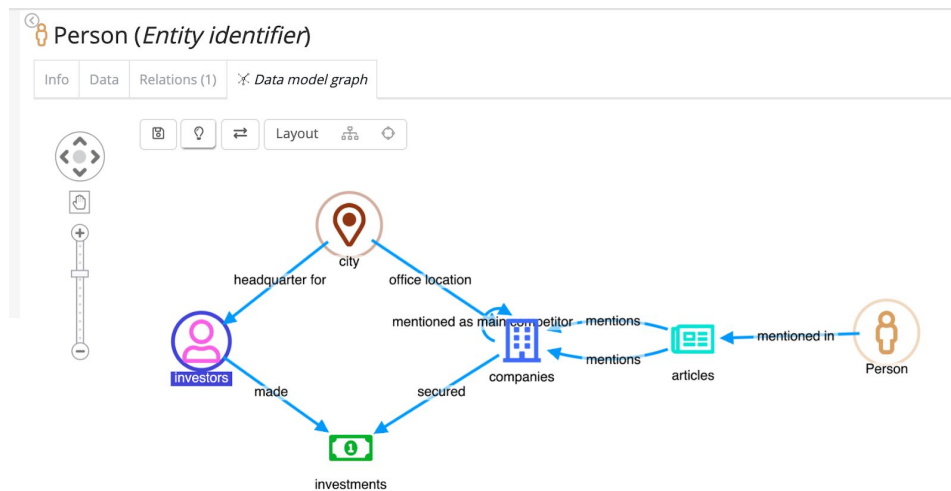
Persons (*Entity identifier*)

Info Data Relations [Data model graph](#)

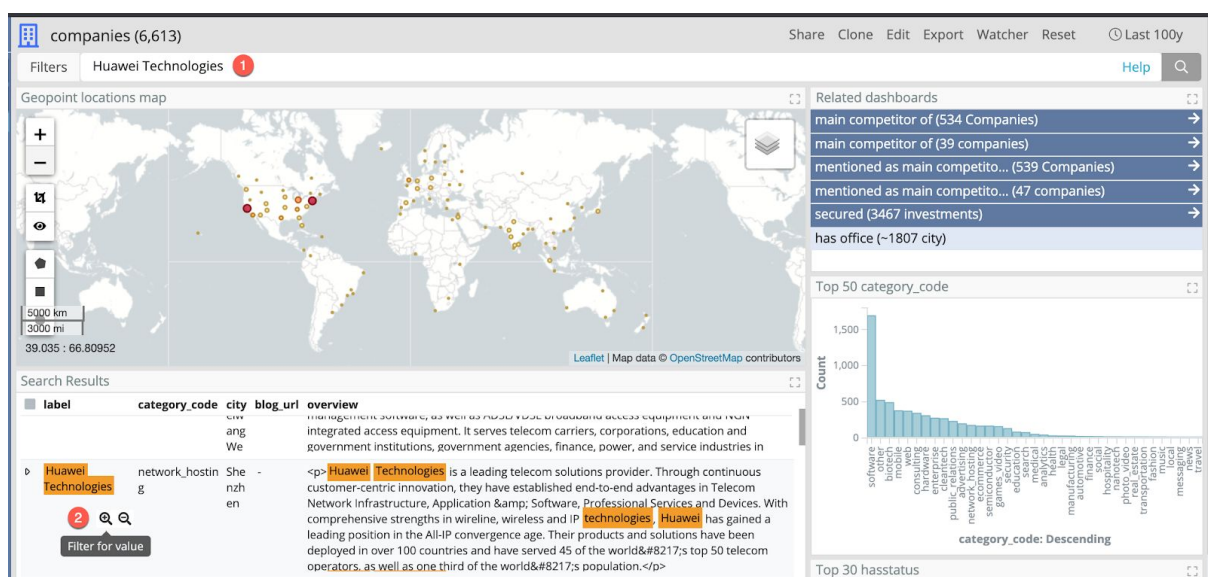
Source Entity	Labels	Target Entity
Persons	→ mentioned in ← mentions	Field: articles siren.nlp.matches.entity/person...

[Add relation +](#) [Relations auto-discovery wizard \(BETA\) >](#)

3. The data model graph should look like this, at this stage:

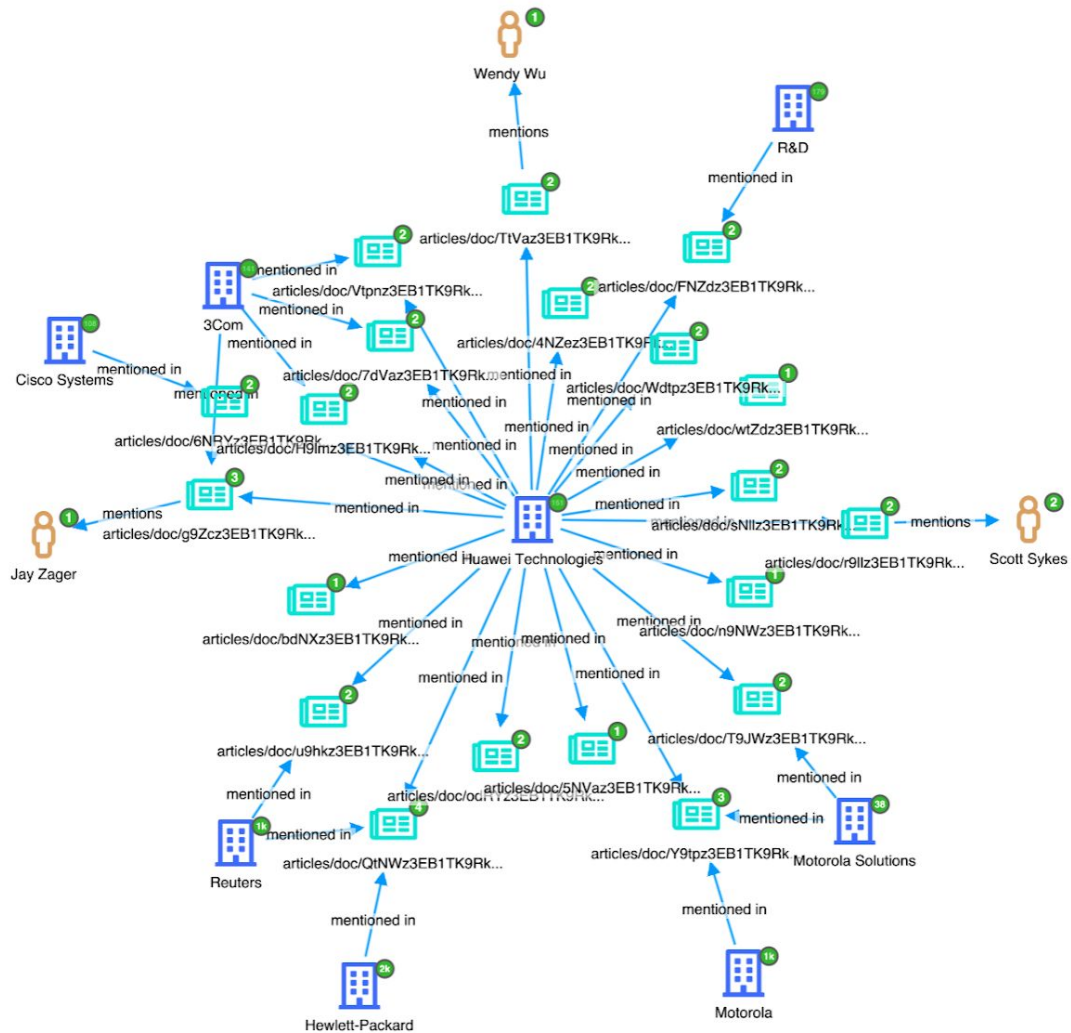


4. Go to the **Companies** dashboard and search for the company “Huawei Technologies” ①.
5. Filter the dashboard for the company with the label name “Huawei Technologies” by using the **Filter for value** button ②. This will filter the dashboard with a single record for this company.



6. Go to the **Graph Browser** and drag and drop the **Companies** dashboard there. This will add a node with the company Huawei Technologies.
7. Select the node, right-click, and select **Expand by relation**.
8. Select **mentioned in (20) articles** and click **Ok**.
9. You will now see the articles that mention this company. Next, select all of the nodes on the graph and right-click → **Expand by relation** and select **mentions (3) Person** and **mentions (32) companies**.

10. This will give us an interesting graph that shows articles mentioning Huawei Technologies and also mention some **other companies**. The graph also shows mentions of **persons** who were named in articles that talk about Huawei Technologies.



Appendix 2: Configuring the graph

You can modify the graph in a number of ways by using the features that are provided in the Graph Browser.

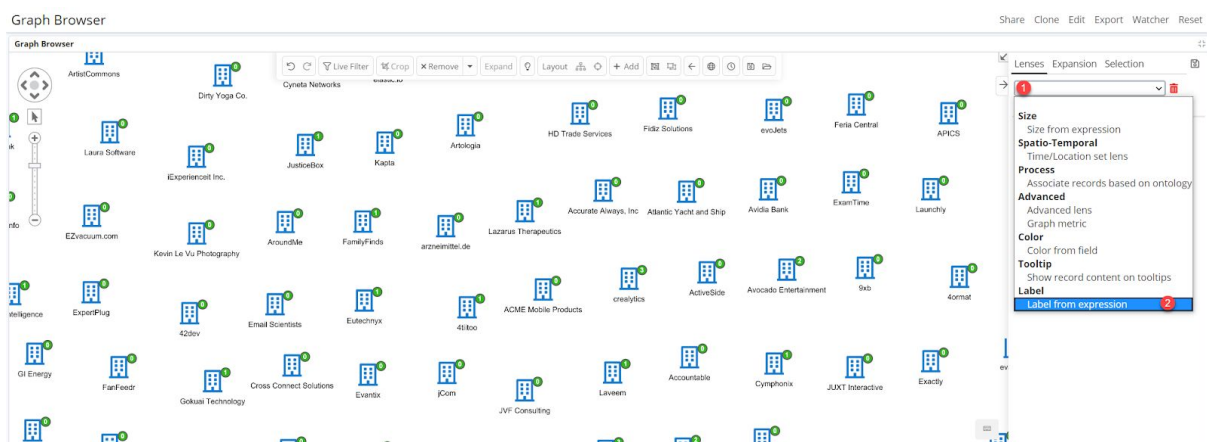
Fixing labels with graph lenses

You can set graph node labels in the data model and then in the graph by using the lens feature.

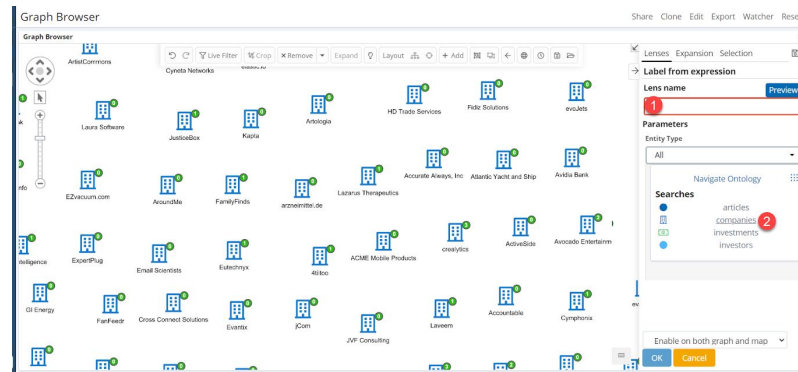
You should always set them in the data model first, but you can then override them in the graph with a lens that is useful for specific analysis. For example, you might first want to see the article's title, but later change your mind and want to see the article's author as a label instead.

To create and activate a lens, complete the following steps:

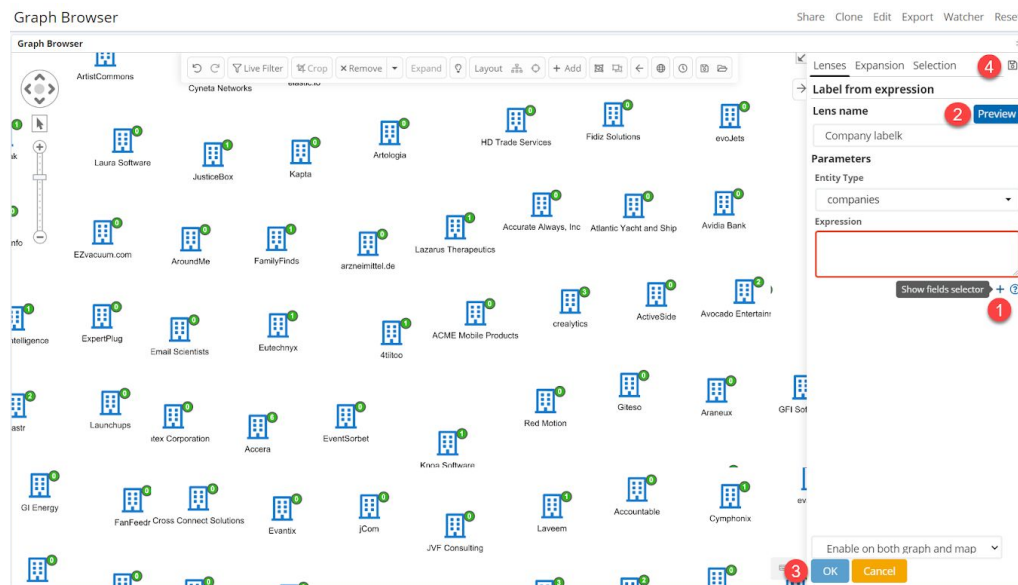
1. Click the **Toggle Sidebar** button ①.
2. In the **Lenses** tab, click **Add a lens** and select **Label from expression** ②.



3. In the Lense name editor, enter “Company labels” ①.



4. In the **Entity Type** drop-down menu, select **companies** from the list of searches ②.
5. To fill the **Expression** field, click the + button ① and select **label** from the dropdown menu. The **Expression** field is populated automatically with `payload["label"]`. Click **Preview** ② and then **OK** ③.
6. Do not forget to save your lenses! Click on the **disk** icon in the top-right of the panel ④.



Using graph cards

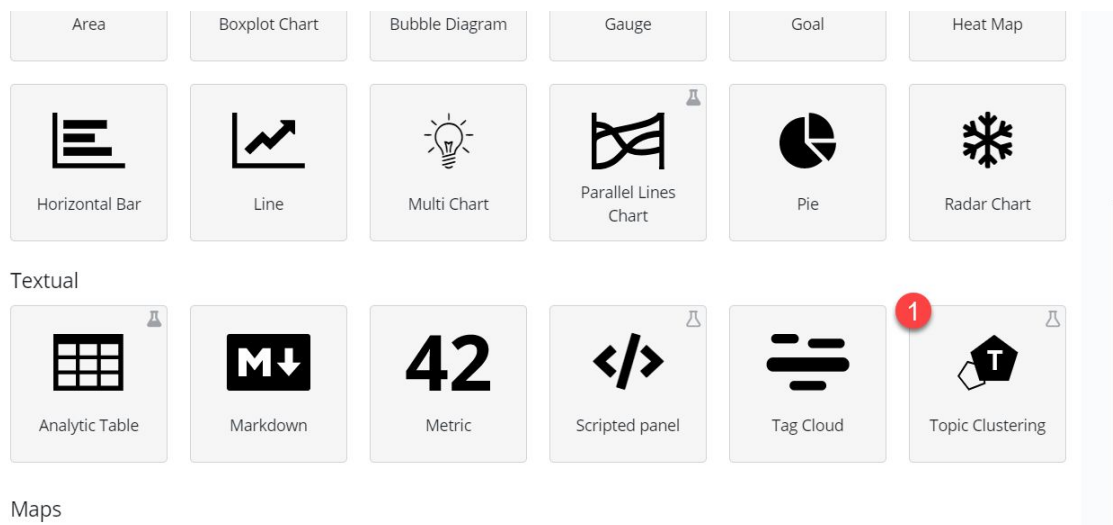
The **Cards** tab in the Graph Browser allows you to create a script to visualize the metadata of the selected nodes, or perform other actions. The cards can be configured by writing a JS-based script in Siren.

For more information, see [Using cards on the graph](#) in the Siren Platform documentation.

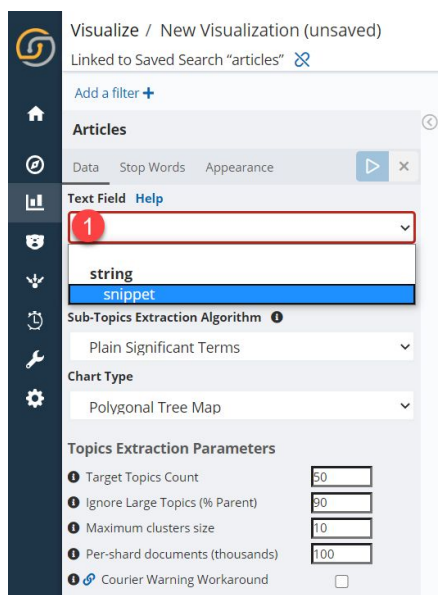
Appendix 3: The Topic Explorer

To better explore the textual content of the articles, go back to the article generated dashboard, click on **Edit** and **Add a new visualization** (pick **articles** as the saved search to be used for this new visualization).

Select the **Topic Clustering** visualization ①.



Select as input the **snippet** field ①, which contains the actual text of the articles.



What is the most discussed topic in our data?



Appendix 4: Creating Entity Identifiers

Siren expands its analysis capabilities by introducing Entity Identifiers(EID).

EIDs are used when entities exist (such as IP addresses, MD5s or things that are identified simply by strings, for example, City) that are present only as *values* in certain fields - in one or more tables - and do not have a “dedicated” table.

EIDs on the graph look like any other node, for example, an IP connecting two log records. They can also be used as ‘intermediate concepts’ when navigating dashboards.

To create an EID, complete the following steps:

1. Go to the **Data Model** page under the **Management** tab and click on **+Create entity identifier**.
2. Enter the name as “city”, choose a suitable icon and color and click **Create**. You will see the entity identifier listed under the Searches on the left (1).
3. Click on the **Relations** tab (2) to create relations for the EID tab. We can create two two relations; one with companies and one with investors, since both of them contain the **city** field.
4. To manually create a relation, click the **Add relation** button, choose the **Target Entity** as **companies**, **Field** as **city** and enter the Labels as in the screenshot below.
5. Do the same for the **investors** entity with the **Field** as **city**.

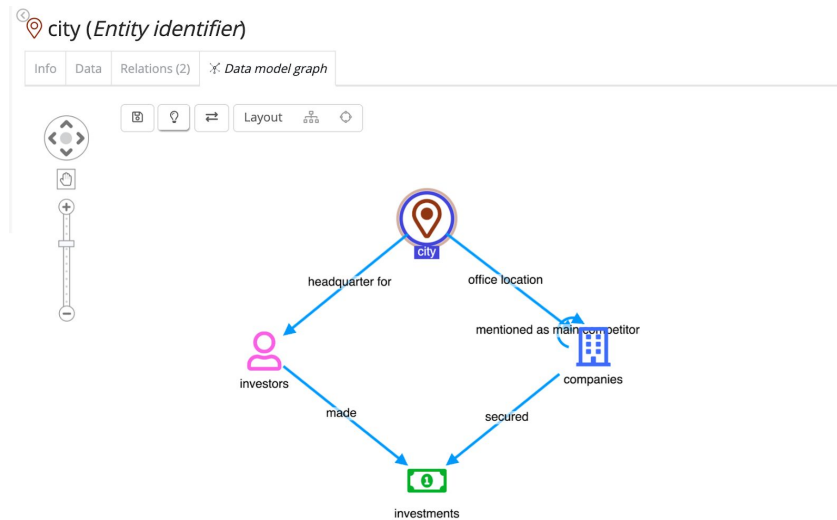
city (Entity identifier)

Info Data Relations Data model graph

Source Entity	Labels	Target Entity
city	→ office location ← has office	companies Field: city
city	→ headquarter for ← headquartered in	investors Field: city

Add relation + Relations auto-discovery wizard (BETA) >

4. Click **Save**.
5. Go to the Data model graph tab to see the EID added as a part of the data model.



Appendix 5: Dashboard 360

The Dashboard 360 feature allows a single dashboard to contain visualizations that are based on different searches, and allows you to perform coherent filtering across all of them. Dashboard 360 uses the relational data model to enhance data analysis capabilities in Siren Platform.

1. Click **Create a new Dashboard** and name it “All Companies”.
2. Click **Add -> Add new Visualization**.
3. Select a **Record Table** visualization and associated it with the **siren-import-home-companies** search.
4. Repeat steps 2-3 to add a **Bar Chart** called “Histogram of funded_year” and associate it with the **siren-import-home-investments** search.

As two searches are involved, we have to use a Dashboard 360. To configure Dashboard 360, complete the following steps:


1. Click **Edit** → **Data Model** and select **Dashboard 360 with filter strategy** (leave the filter strategy as **Left filter join**) ①.
2. Click **Add Main Search** ②.

Dashboard Data Model ?

☐ Dashboard does not represent a search

☐ Dashboard is about a search

☒ Dashboard 360 with filter strategy Left filter join




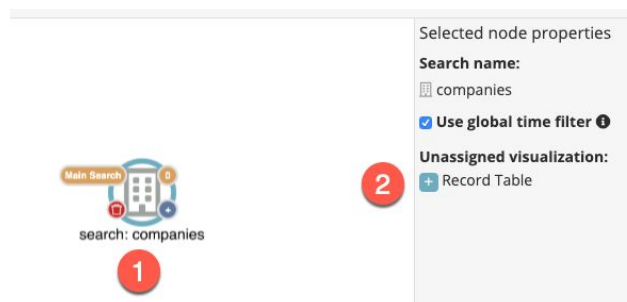
This opens a dialog that enables you to select the appropriate search.

3. Select **siren-import-home-companies** as the main search and click **Done**.

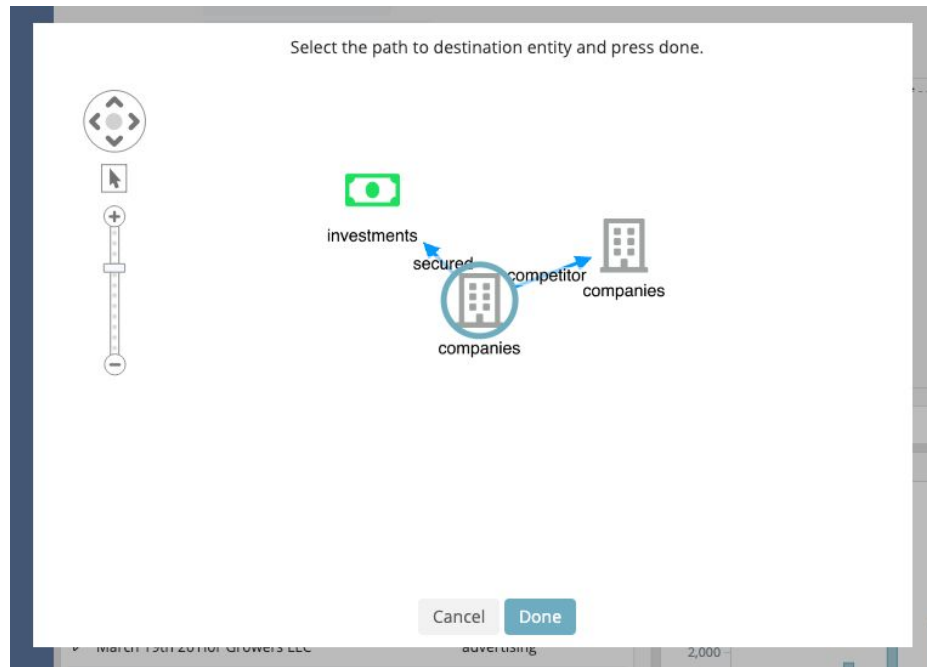
By default, the side panel shows the list of Dashboard 360-compatible visualizations and whether or not they are assigned to a search node. In this case, neither of the visualizations created earlier are assigned yet.




4. To assign a visualization, click the **companies** node to select it ①, then click the  button beside the required visualization (Record Table) ②.




5. You can now build the dashboard data model by defining how this search is to be relationally connected with other searches (that is, which path in the relational data model to take). To do this, click the \oplus button located on the lower-right of the Main Search node. This opens the following dialog which shows all the possible relations that come from the chosen search, as well as all possible searches that can be reached through those relations.



6. Select Investments, via the only possible relation (secured), then click **Done**.
7. You can now see how the dashboard data model looks at this point, with the companies node selected.

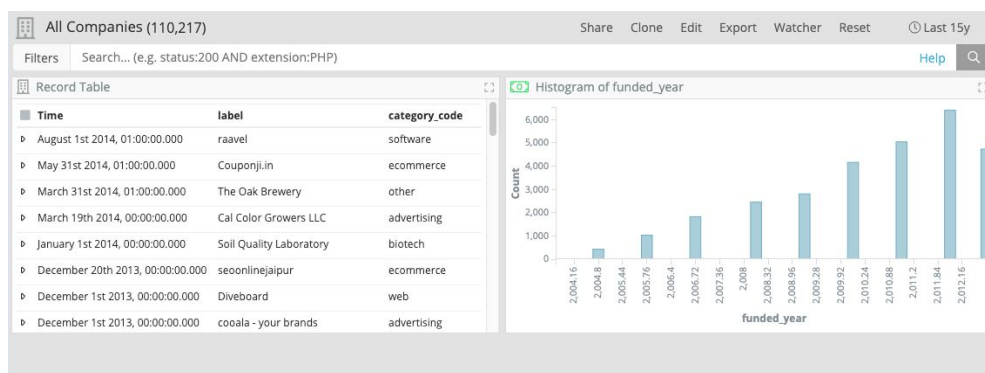


Select the **investments** node and you will see that it has an unassigned visualization.



8. Assign this visualization by clicking the ⊕ button, then save the dashboard.

Let's see what the dashboard looks like now that the visualizations have been assigned. Note the icons on the top-left of each visualization's container, indicating they have been assigned to the appropriate searches.



Filtering with Dashboard 360

The really significant achievement here with Dashboard 360 is that you can perform coherent filtering across the dashboard's visualizations.

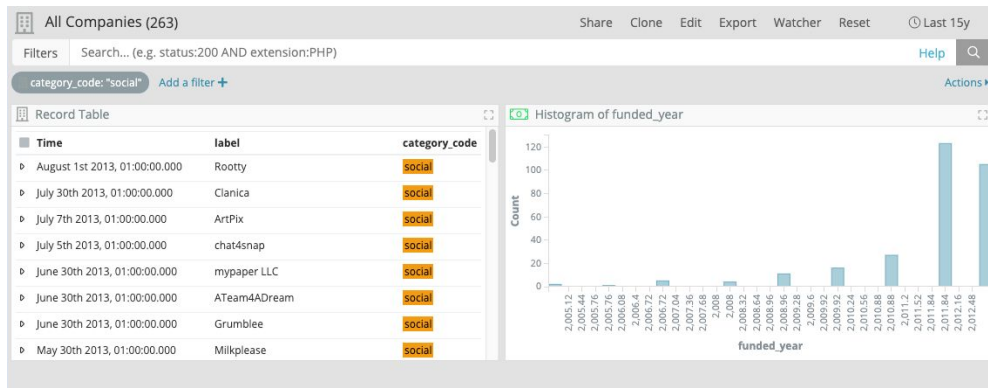
Let's say that you want to see the pattern of yearly funding across a particular industry segment. For example, when did funding for social media really take off?

To do this, filter the Record Table visualization on the appropriate category code, **social**.

August 1st 2013, 01:00:00.000	Raploja	ecommerce
August 1st 2013, 01:00:00.000	Rootty	social
August 1st 2013, 01:00:00.000	Epic Business	const

Filter for value

The Record Table now lists all of the companies in the social category. But, more importantly, the other visualization displays the yearly funding trend; this showed a huge increase in the number of investments between 2010 and 2011.



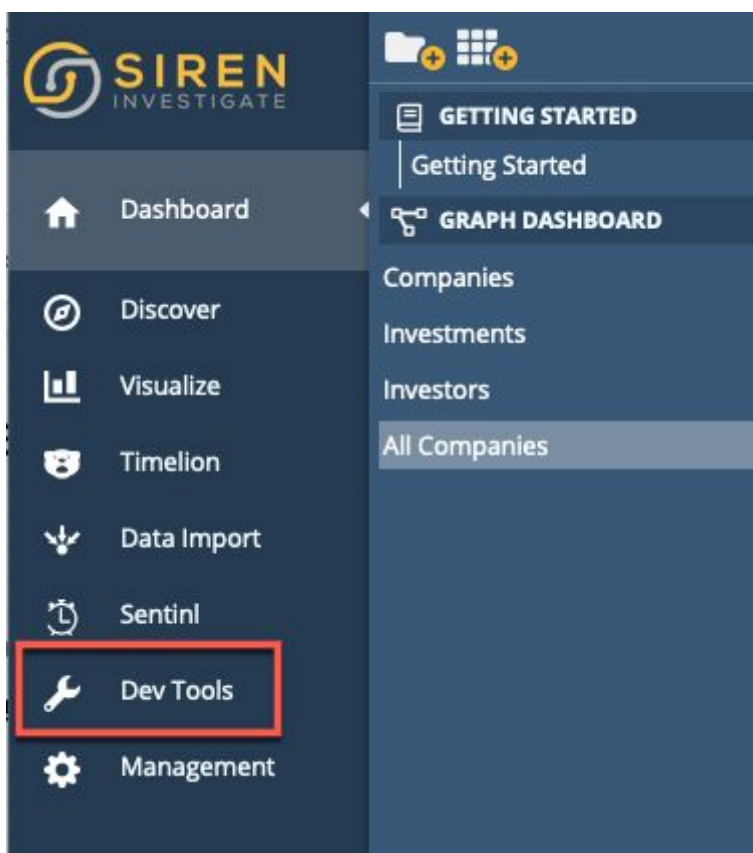
You can also filter by clicking on the segments in the **Funded Year** histogram, this will coherently filter the results in the Record Table.

Appendix 6: Deleting data in Siren and changing a table schema

Deleting data in Siren Platform is intentionally difficult, due to the damage that it could cause in environments where the data is in very large streams, such as in large Elasticsearch installations.

However, you can delete an index by using the Dev Tools console.

1. Go to the **Dev Tools** app.



2. In the **Console**, type “delete” and the name of the index, for example “siren-” in the editor pane.
3. Click the play icon ① and you see the acknowledgment in the response pane.

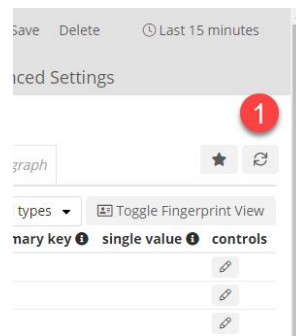
Dev Tools

Console Translate Gremlin Console



IMPORTANT: If you delete the index and the **companies** dashboard still exists, the dashboard displays errors. You can either reinstate the companies index by uploading the data again, or you can delete the dashboard.

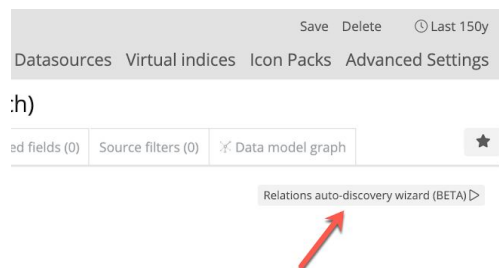
If the data structure has changed, for example, if there are more fields or fields of a different type, then you must refresh the 'Index Pattern Search' field list by clicking the **Refresh** button ①.



Appendix 7: Auto-discovery of relations

This feature automatically detects and creates relations between index pattern searches. To use this option, complete the following steps:

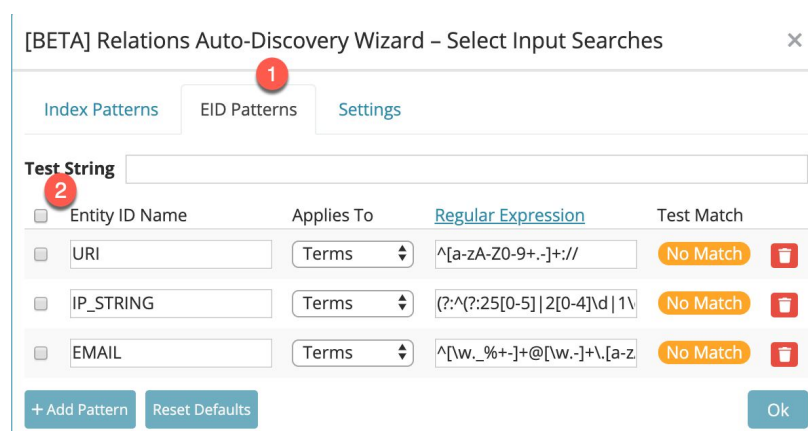
1. Go to **Management** → **Data Model** and click on the **Relations** tab.
2. Click on **Relations auto-discovery wizard**.



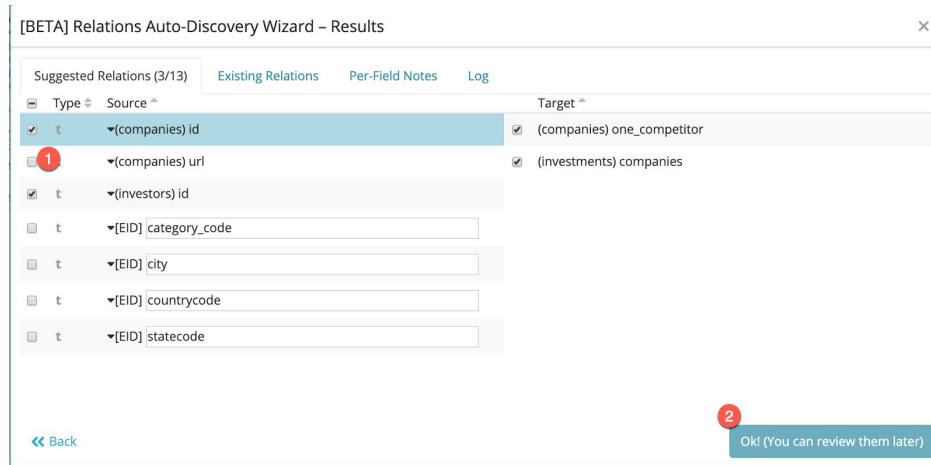
3. The wizard automatically selects all of the Index Patterns and also adds **EID Patterns**.

NOTE: EIDs are created to link certain shared identifiers that belong to certain entities; for example, a URL, an email address, or a city. However, there is no corresponding index in which these entities are the primary key. They are simply used as *values* in fields across different indices and hence cannot be directly used to create a relation. EID will be discussed in more detail later in this tutorial.

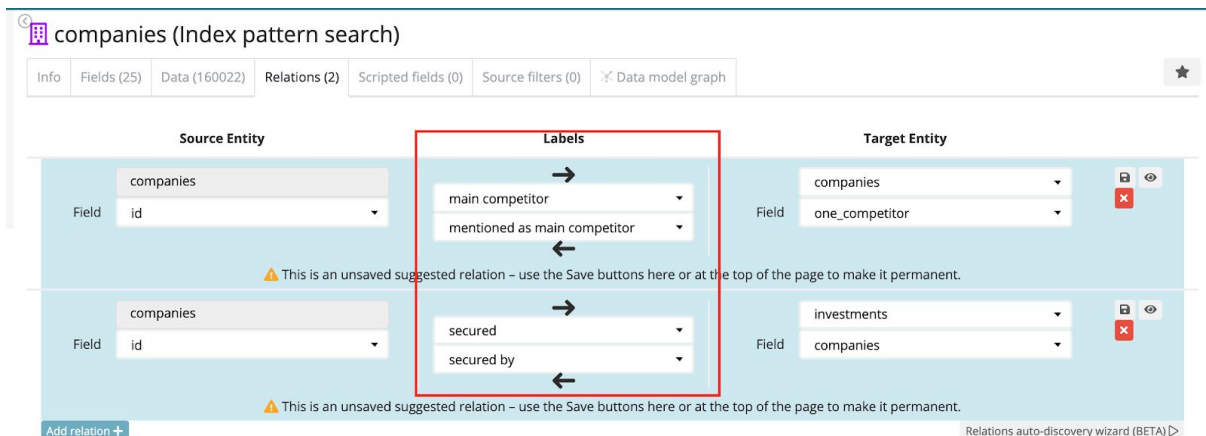
4. Click on the **EID Patterns** tab (1) and deselect all of the EID name fields (2), since we're only interested in creating relations between index patterns for now. Click **Ok**.



5. You can also deselect the additional fields for an index such as **(companies) url** (1) since **(companies) id** is already selected for the companies index. This will simplify the data model for us. Click **Ok** (2).

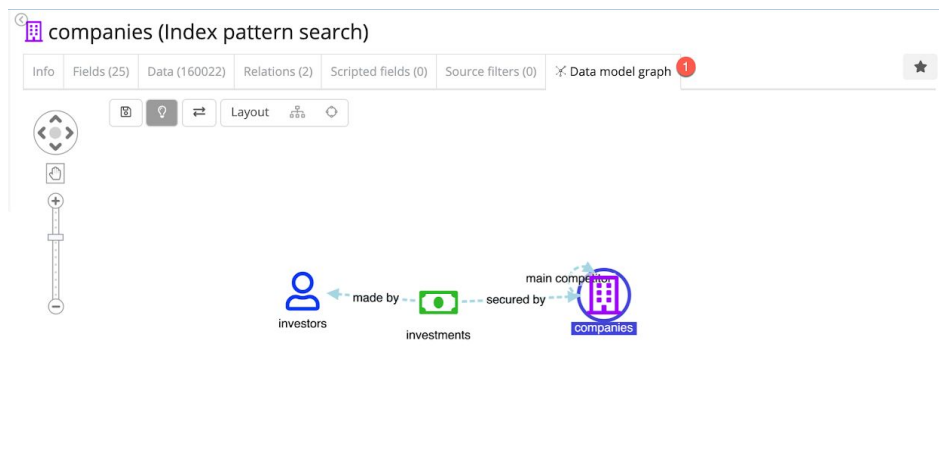


6. The relations are automatically generated in a draft state (highlighted in blue). You can review the relations that are created for all of the indexes.
7. Change the **Labels** for the relations to make the relations clear, as shown in the screenshots below.



8. Click the **Data model graph** tab to see the data model in a visual representation.
9. Click **Save All** to save the relations.

In the screenshot below, you can see the relational data model that is generated for the entities, with the relations **Investors** → *made* **Investments** → *secured by* **Companies** → *with the main competitor of other* **Companies**.



Legal notices

Siren Platform™ is a trademark of Sindice Ltd. trading as Siren, with offices in GTC, Mervue Business Park, Galway, Republic of Ireland.

Elasticsearch™ is a trademark of Elasticsearch B.V., registered in the U.S.A. and in other countries.