

DEPLOYMENT GUIDE

Elastic Stack + Infoblox B1TD & NIOS DNSTAP Logging Integration

March 2021



Table of Contents

Introduction	2
Requirements	2
For Integrating BloxOne Threat Defense Data	2
For Integrating NIOS dnstap Logs	2
Tested Hardware & Software	2
Install Elastic Stack	3
Deployment Instructions: B1TD	3
CSP API Key Retrieval	3
Python Configuration	4
Logstash Configuration for TD	7
Kibana Data Discovery	9
Kibana Data Visualization	14
Deployment Instructions: NIOS dnstap Logging	21
NIOS Configuration	21
dnstap-receiver Configuration	22
Logstash Configuration for dnstap	24
Run & Test Configuration	26

Introduction

Elastic Stack, formerly known as the ELK Stack, is a popular suite of tools that provides advanced logging, storing, searching and visualization functionality to data of many types from any source. Elasticsearch, Logstash, Kibana, and newcomer Beats work together to make up the core products of Elastic Stack. Elasticsearch handles search and storage of data, Logstash is the pipeline for retrieving data to send to Elasticsearch, and Kibana provides the web browser user interface used to visualize and query this data. Elastic Stack is available as a free, open source local download, but it also provides a paid-for cloud solution. We will be working with the open source version.

This deployment guide is two-fold. As well as providing a detailed introduction into Elastic, it contains deployment guides for integrating both BloxOne Threat Defense data and NIOS dnstap logging data.

Drastically enhance the ability to analyze your network by integrating Elastic Stack's powerful data exploration tools with Infoblox's extensive security and query/response data.

Requirements

For Integrating BloxOne Threat Defense Data

The following items are required to incorporate Infoblox BloxOne Threat Defense DNS security data into Elastic Stack:

- Access to an Infoblox BloxOne Threat Defense subscription
- Access to an Elastic Stack instance
 - Composed of Elasticsearch, Logstash and Kibana

For Integrating NIOS dnstap Logs

The following items are required to incorporate NIOS dnstap for high-performance query logging into Elastic Stack:

- A NIOS Grid with dnstap enabled
- Access to an Elastic Stack instance
 - Composed of Elasticsearch, Logstash and Kibana
- A dnstap receiver to read and translate dnstap logs from NIOS

Tested Hardware & Software

- Windows 10 VM
- Ubuntu 18.04 VM
 - Elastic Stack version 7.9.2 installed
 - Elastic Stack composed of Elasticsearch, Logstash and Kibana
 - Python3 installed
 - (Optionally) [dnstap-receiver module](#) for Python installed
- (Optionally) IB-FLEX NIOS VM with DNS Cache Acceleration and dnstap enabled
 - Running NIOS 8.5.2

Install Elastic Stack

There are several ways to install Elastic Stack. You may prefer different package formats or operating systems depending on your needs and preferences. Find more information and further links for installing each Elastic Stack component at <https://www.elastic.co/guide/en/elastic-stack/current/installing-elastic-stack.html>.

A comprehensive one-all guide for installing and configuring Elastic Stack and its dependencies on Ubuntu 18.04/20.04 can be found at <https://phoenixnap.com/kb/how-to-install-elk-stack-on-ubuntu>.

Deployment Instructions: B1TD

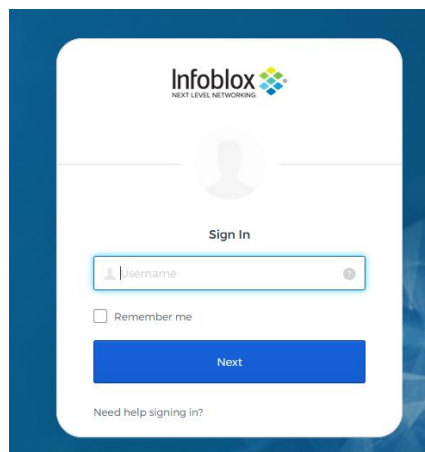
The following instructions provide an intro into the Elastic Stack as well as the steps required to integrate BloxOne Threat Defense DNS Security data with Elastic.

CSP API Key Retrieval

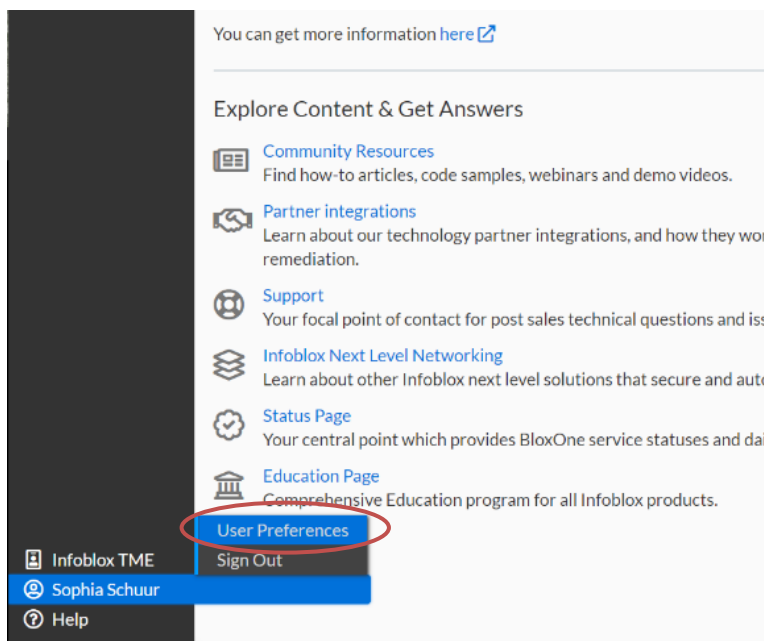
You will need a BloxOne Threat Defense API key to pull the DNS data. You can access this key through the Cloud Services Portal (CSP). API keys are unique identifiers found in many applications to both identify the application making the API calls and verify the application making the calls has access to do so.

To access your API key:


1. Log into the CSP at <https://csp.infoblox.com>.



2. Upon logging in, hover over your username in the bottom -left corner of the CSP and select **User Preferences**.



3. A popup will appear. Click **Copy** to copy your API key to your clipboard. Copy it somewhere you can easily access and copy from later, such as Notepad. This will be the key you copy into the Python script later.



Python Configuration

We will be using a Python script to gather the most recent ten minutes of DNS event data from Infoblox's BloxOne ThreatDefense REST API and write it into `json` log files. To ensure the data is always recent and updated, we will tell Ubuntu to run this script every ten minutes. Later we will configure Logstash to read the `json` and send it to Kibana to be visualized.

Let's create the Python script. You can save this script anywhere easily accessible to you, such as Documents or the Desktop. *Note: For this demo we will be saving it in `/home/<username>/dataconnector`.*

1. Access the machine where your Logstash instance is installed. *Note: For this demo Elastic Stack was installed on Ubuntu 18.04.*
2. Open a terminal.
3. **Python3** must be installed for the script to work properly. If it is not already installed, install it with:

```
sudo apt install python3.8
```

4. Navigate to `/home/<username>/dataconnector`:

```
cd /home/infoblox/dataconnector
```

5. Create a new Python file:

```
touch cspscript.py
```

6. Open the file with `gedit` for editing:

```
gedit cspscript.py
```

7. Copy and paste the following into the file. **Careful to note Python's spacing and tabbing syntax. Indent nested Python newlines with four spaces.** Replace the text `<YOUR API KEY HERE>` with the BloxOne TD API key acquired in the [CSP API Key Retrieval](#) section of this document. Save and close the file when finished. To ensure your formatting is correct, you can also download the script [here](#) on InfobloxOpen's Github repo.

```
# -*- coding: utf-8 -*-

import os
import datetime
import calendar
import requests
import json
import time

now = datetime.datetime.utcnow()
ten_minutes_ago = datetime.datetime.utcnow() - datetime.timedelta(minutes = 10)
filename = f"{ten_minutes_ago.strftime('%Y%m%d_%H%M%S')}_{now.strftime('%H%M%S')}"

sif_now = calendar.timegm(now.timetuple())
sif_last_hour = calendar.timegm(ten_minutes_ago.timetuple())

url = f"https://csp.infoblox.com/api/dnsdata/v1/dns_event?t0={sif_last_hour}&format=json&t1={sif_now}&source=rpz"

time.sleep(120 )
payload = {}
headers = {
    'Authorization': 'Token <YOUR API KEY HERE>'
}
response = requests.request("GET", url, headers=headers, data = payload)

path = "/tmp/rpz"
if not os.path.exists(path):
    os.makedirs(path) ← Watch the spacing here!

completeName = os.path.join(path, filename+".json")
data = json.loads(response.content)
write_data = json.dumps(data)

fh = open(completeName, 'w')
fh.write(write_data)
fh.close()
```

Let's make sure the script is working. Run:

```
/usr/bin/python3 /home/<username>/dataconnector/cspscript.py &
```

You should see an output like this:

```
infoblox@infoblox-virtual-machine:~$ /usr/bin/python3 /home/infoblox/dataconnector/cspscript.py &
[7] 12705
[6] Done /usr/bin/python3 /home/infoblox/dataconnector/cspscript.py
```

The software utility cron is a time-based job scheduler in Unix-like operating systems. Let's configure cron to run the Python script every ten minutes.

1. Open a terminal.
2. Open a crontab file for editing:

```
crontab -e
```

3. Choose your preferred editor if prompted.

```
infoblox@infoblox-virtual-machine:~/sdataconnector$ crontab -e
no crontab for infoblox - using an empty one

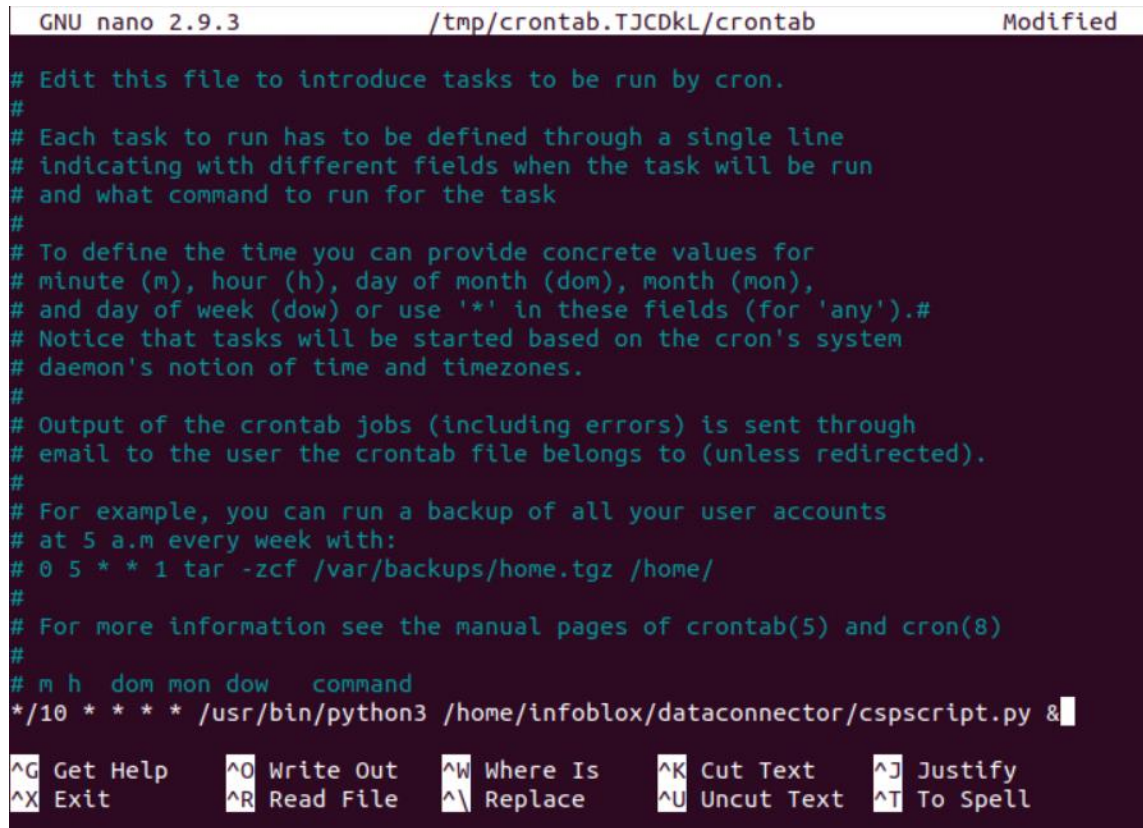
Select an editor. To change later, run 'select-editor'.
 1. /bin/nano      <---- easiest
 2. /usr/bin/vim.tiny
 3. /bin/ed

Choose 1-3 [1]: 1
Use "fg" to return to nano.

[1]+  Stopped crontab -e
```

4. Insert the following line into the file as shown below. Save and exit the editor.

```
*/10 * * * * /usr/bin/python3 /home/<username>/dataconnector/cspscript.py &
```



```
GNU nano 2.9.3 /tmp/crontab.TJCDkL/crontab Modified
# Edit this file to introduce tasks to be run by cron.
#
# Each task to run has to be defined through a single line
# indicating with different fields when the task will be run
# and what command to run for the task
#
# To define the time you can provide concrete values for
# minute (m), hour (h), day of month (dom), month (mon),
# and day of week (dow) or use '*' in these fields (for 'any').#
# Notice that tasks will be started based on the cron's system
# daemon's notion of time and timezones.
#
# Output of the crontab jobs (including errors) is sent through
# email to the user the crontab file belongs to (unless redirected).
#
# For example, you can run a backup of all your user accounts
# at 5 a.m every week with:
# 0 5 * * 1 tar -zcf /var/backups/home.tgz /home/
#
# For more information see the manual pages of crontab(5) and cron(8)
#
# m h dom mon dow  command
*/10 * * * * /usr/bin/python3 /home/infoblox/dataconnector/cspscript.py &
```

Logstash Configuration for TD

Logstash is a highly customizable part of Elastic Stack that retrieves data. It can be configured to collect data from many different sources, such as log files, REST API requests, and more, to be sent to Elasticsearch and later visualized in Kibana. Hundreds of plugins are available to expand its functionality, and many are included with the software at installation. We will be using the input plugin `file` to read the `json` logs generated by the Python script. A complete list of available plugins and links to their documentation can be found at https://www.elastic.co/support/matrix#matrix_logstash_plugins.

Logstash configuration is governed by special configuration files. Where to retrieve data, how to filter it, and where to output it are configured by these files. For this demo, Elastic Stack was installed on Ubuntu 18.04 via `apt-get`, so these files are set by default to live in the `/etc/logstash/conf.d` directory. Your directories may be different depending on how Elastic Stack was installed. More information about the Logstash directory layout can be found at <https://www.elastic.co/guide/en/logstash/current/dir-layout.html>.

Let's configure Logstash to grab the DNS security data found in the `json` files generated by the Python script.

1. Access the machine where your Logstash instance is installed. *Note: For this demo Elastic Stack was installed on Ubuntu 18.04.*
2. Open a terminal.

3. Navigate to where your Logstash configuration (.conf) files are located. In this demonstrative environment, these files are located in /etc/logstash/conf.d. Input the following command to navigate to the correct directory:

```
cd /etc/logstash/conf.d
```

4. Create a new file called csp-dns-events.conf:

```
sudo touch csp-dns-events.conf
```

5. Open the file with gedit for editing:

```
sudo gedit csp-dns-events.conf
```

6. Copy and paste the following into the file. Save and close the file when finished.

```
input {
  file {
    path => "/tmp/rpz/*"
    codec => "json"
    mode => "read"
    syncedb_path => "/dev/null"
  }
}

filter {
  split {
    field => ["result"]
    terminator => ","
  }
  mutate {
    remove_field => ["status_code"]
  }
}

output {
  elasticsearch {
    hosts => ["localhost:9200"]
    index => "csp-dns-events"
  }
}
```

Let's get a breakdown of what is happening in this code.

- a. **Input:** Here is where we read the json files created by the Python script. We use the input plugin [file](#).

- b. **Filter:** This is where we [split](#) every record found in the `json` files into individual hits in Kibana. By doing so, we can directly search and organize by any field returned by the GET request in the Python script. The returned body of the request is placed in one field called `result` with each record terminated by a comma, so we tell that to Logstash. Using the [mutate](#) plugin we remove the extra `status_code` field, since it creates unnecessary clunky data.
- c. **Output:** Send the data to Elasticsearch and give this index a name. This is the name that will appear in Kibana when creating a new index.

Logstash config files follow a specific schema. More information on the structure of config files can be found at <https://www.elastic.co/guide/en/logstash/current/configuration-file-structure.html>

7. Navigate to your home directory for Logstash. For this demo, this is `/usr/share/logstash/`. Input the following command to navigate to the correct directory:

```
cd /usr/share/logstash
```

8. Run Logstash with your new configuration:

```
sudo bin/logstash -f /etc/logstash/conf.d/csp-dns-events.conf
```

Allow several minutes of processing. The console will inform you if there are any syntax errors with your config file.

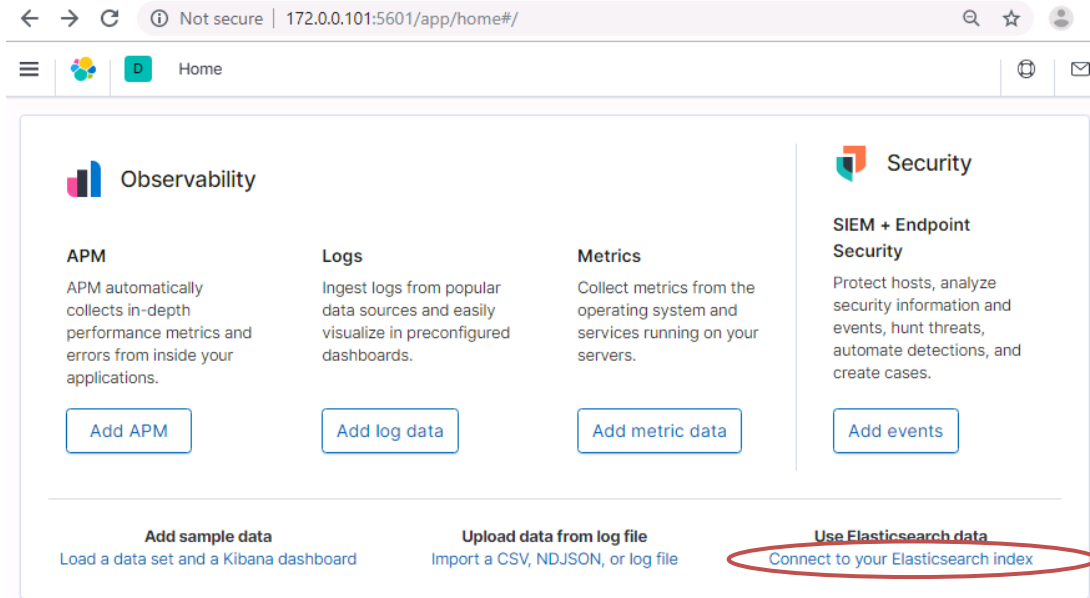
Alternatively, you can simply restart the Logstash service, but the console will not warn you of any errors with your config file:

```
sudo systemctl restart logstash
```

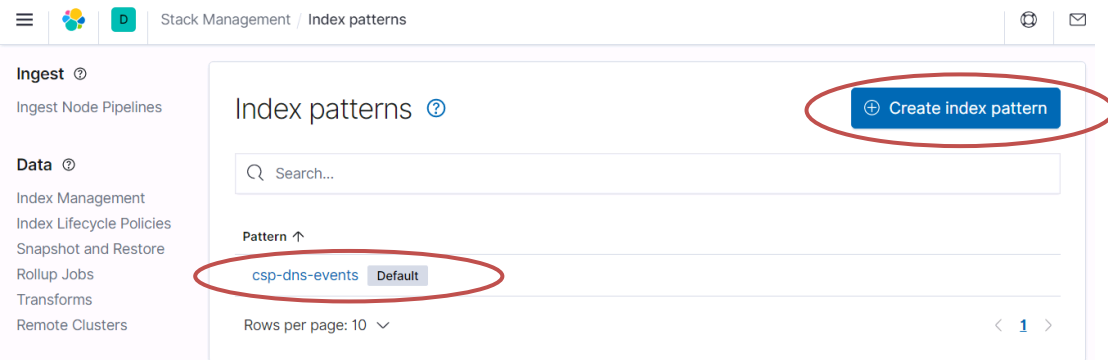
Kibana Data Discovery

Kibana is the visualization part of the Elastic Stack. It provides a web-based user interface for viewing and charting the data stored in Elasticsearch. Using Index Patterns, we can map Kibana with the data that our Logstash configuration is outputting to Elasticsearch.

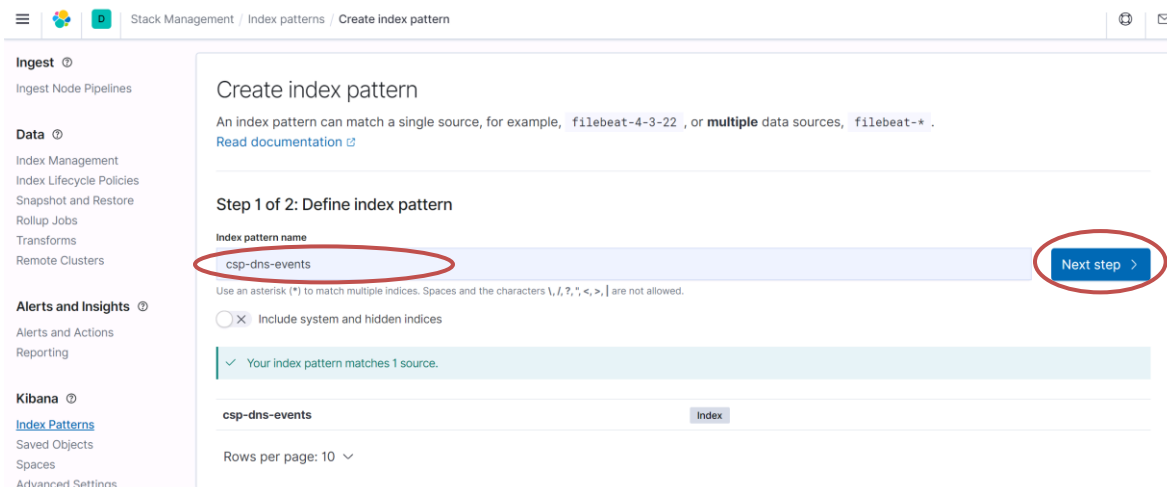
1. Access your Kibana instance. *Note: If desired, you must configure Kibana to allow remote access, such as from a secondary Windows machine. Find instructions [here](#).*
2. From the home page, click **Connect to your Elasticsearch index**.



3. Your configuration will appear here as an available index pattern. Click **Create index pattern**.



4. Name the index pattern **csp-dns-events**. Then click **Next step**.



5. In the Time field, select **@timestamp**. Then click **Create index pattern**.

Create index pattern

An index pattern can match a single source, for example, `filebeat-4-3-22`, or **multiple** data sources, `filebeat-*`.
[Read documentation](#)

Step 2 of 2: Configure settings

csp-dns-events*


Select a primary time field for use with the global time filter.




Time field Refresh

@timestamp

> [Show advanced options](#)


< Back **Create index pattern**

6. Click the  menu icon in the topbar. Select **Discover**.

   Stack Management

Home

Recently viewed >

 Kibana

Discover

Dashboard

Here you will see all the DNS security records retrieved from the CSP.

The screenshot shows the Discover interface with the following components:

- Search Bar:** Search, KQL, Date range: Oct 27, 2020 @ 18:14:30.00 → Oct 27, 2020 @ 18:15:00.00, Refresh button.
- Filters:** + Add filter, Filter by type (0).
- Selected fields:** _source.
- Available fields:** _id, _index, _score, _type, @timestamp, @version, host, path, result.category, result.confidence, result.country, result.device, result.dhcp_fingerprint, result.event_time, result.feed_name.
- Chart:** 52 hits, Oct 27, 2020 @ 18:14:30.000 - Oct 27, 2020 @ 18:15:00.000, Auto.
- Table:**

Time	_source
> Oct 27, 2020 @ 18:14:45.503	host: infoblox-virtual-machine path: /tmp/rpz/20201028_010443_011443.json @timestamp: Oct 27, 2020 @ 18:14:45.503 @version: 1 result.feed_type: IP Address result.rcode: PASSTHRU result.country: UA result.qtype: A result.event_time: Oct 27, 2020 @ 18:13:49.000 result.os_version: result.feed_name: EECN_IP result.qip: 192.168.0.123 result.category: result.qname: a.ntps.org result.mac_address: result.rdata: 185.126.112.98 result.network: smakousky-NIOS-DFP (DFP) result.dhcp_fingerprint: result.device: 192.168.0.123 result.policy_name: smakousky-security-policy
> Oct 27, 2020 @ 18:14:45.503	host: infoblox-virtual-machine path: /tmp/rpz/20201028_010443_011443.json @timestamp: Oct 27, 2020 @ 18:14:45.503 @version: 1 result.feed_type: FQDN result.rcode: PASSTHRU result.country: US result.qtype: A result.event_time: Oct 27, 2020 @ 18:13:34.000 result.os_version: result.feed_name: smithj whitelist result.qip: 192.168.0.2 result.category: result.qname: occ-0-465-472.1.nflxso.net result.mac_address: 00:0c:29:61:25:a4 result.rdata: 198.38.111.131 result.network: smithj-DFP-05 (DFP) result.dhcp_fingerprint: result.device: 192.168.0.2 result.policy_name: smithj policy
> Oct 27, 2020 @ 18:14:45.503	host: infoblox-virtual-machine path: /tmp/rpz/20201028_010443_011443.json @timestamp: Oct 27, 2020 @ 18:14:45.503 @version: 1 result.feed_type: FQDN result.rcode: PASSTHRU result.country: US result.qtype: A result.event_time: Oct 27, 2020 @

7. You can perform extensive searching, displaying and filtering here. Add some Available fields in the left panel to view field totals for this dataset and organize your hits.

The screenshot shows the Discover interface with the following components:

- Search Bar:** Search, KQL, Date range: Oct 27, 2020 @ 18:14:30.000 → Oct 27, 2020 @ 18:15:00.000, Auto.
- Filters:** + Add filter, Filter by type (0).
- Selected fields:** result.country, result.tclass, result.property.
- Available fields:** (empty).
- Chart:** 52 hits, Oct 27, 2020 @ 18:14:30.000 - Oct 27, 2020 @ 18:15:00.000, Auto.
- Table:**

Time	result.country	result.property	result.tclass
> Oct 27, 2020 @ 18:14:45.503	UA	CountryBlock	Policy
> Oct 27, 2020 @ 18:14:45.503	US	smithj whitelist	CUSTOM
> Oct 27, 2020 @ 18:14:45.503	US	smithj whitelist	CUSTOM
> Oct 27, 2020 @ 18:14:45.503	US	smithj whitelist	CUSTOM
> Oct 27, 2020 @ 18:14:45.503	US	smithj whitelist	CUSTOM
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	smithj whitelist	CUSTOM
> Oct 27, 2020 @ 18:14:45.503	DE	Mobile	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	smithj whitelist	CUSTOM
> Oct 27, 2020 @ 18:14:45.503	US	Generic	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Generic	MalwareC2

8. Use the top search bar to return hits with specified field values. Try searching for “**result.tproperty : Spyware**”. You can save searches, open searches, apply filters and more.

New Save Open Share Inspect

result.tproperty : Spyware KQL Oct 27, 2020 @ 18:14:30.00 → Oct 27, 2020 @ 18:14:30.00

+ Add filter

csp-dns-events

Search field names

Filter by type 0

Selected fields

- result.country
Top 5 values in 7 / 7 records
US 100.0%
- result.tclass
Top 5 values in 7 / 7 records
MalwareC2 100.0%
- result.tproperty
Top 5 values in 7 / 7 records
Spyware 100.0%

Available fields


- _id
- _index
- _score
- _type
- @timestamp

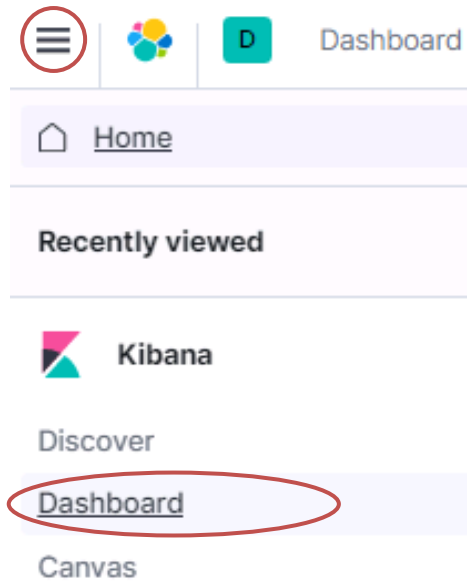
7 hits
Oct 27, 2020 @ 18:14:30.000 - Oct 27, 2020 @ 18:15:00.000 Auto

Time	result.country	result.tproperty	result.tclass
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2
> Oct 27, 2020 @ 18:14:45.503	US	Spyware	MalwareC2

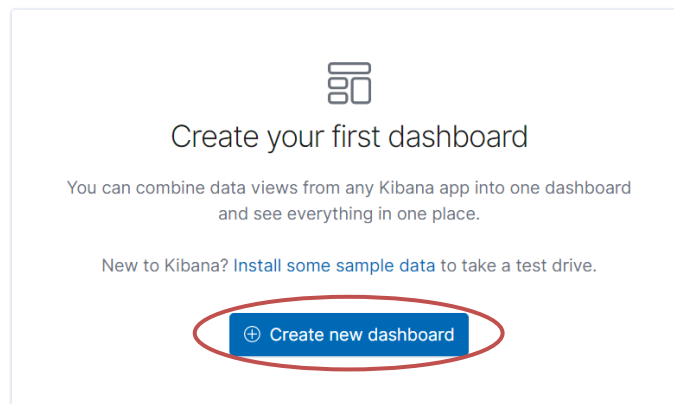
Kibana Data Visualization

Kibana offers many ways of charting data. Let's build a pie chart based on the field **tclass**.

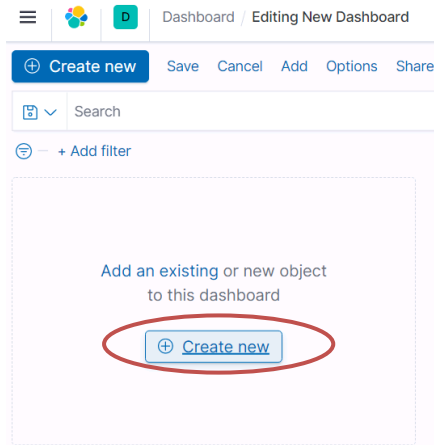
1. Click the  menu icon in the topbar. Select **Dashboard**.



2. Click **Create new dashboard**. We'll add the chart to this dashboard. Many objects, such as charts, can live on a dashboard for convenient access and visualization.

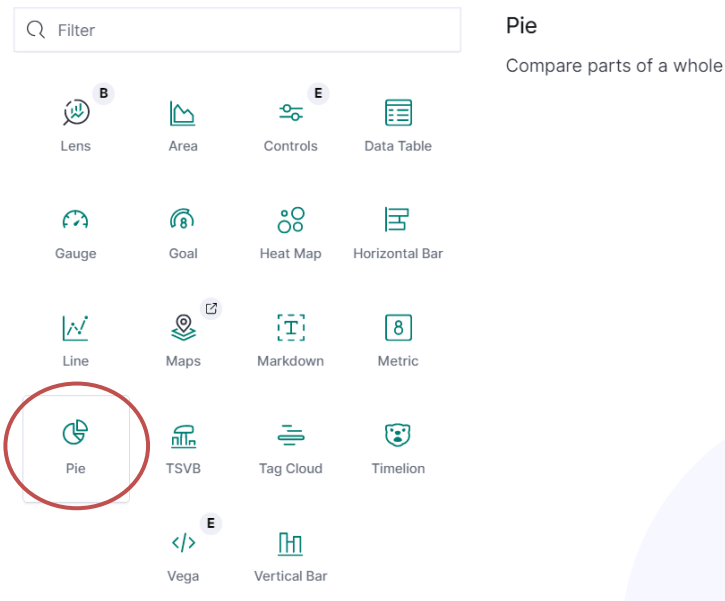


3. Click **Create new**.



4. Select **Pie**.

New Visualization



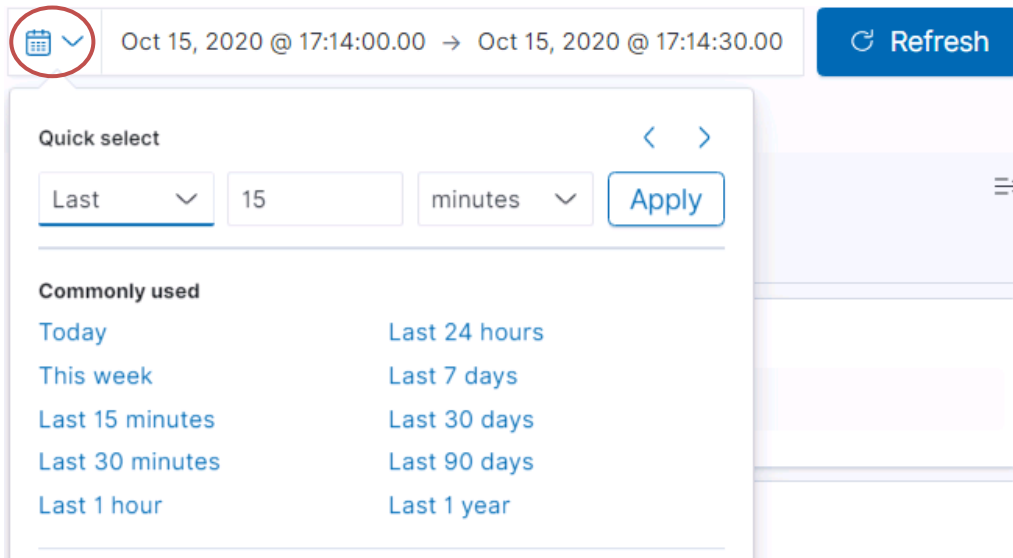
5. Select your index **csp-dns-events**.

New Pie / Choose a source

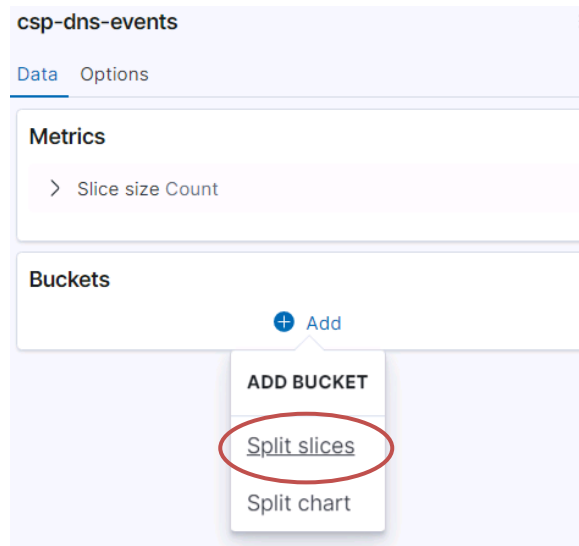


6. You will see an empty pie chart. Let's build it up.

- a. (Optional) If desired, narrow your dataset by clicking on the  calendar dropdown.





- b. Under **Buckets**, click **Add** → **Split slices**.




c. Define the Bucket.


Buckets

Split slices  


Aggregation [Terms help](#)

Terms 


Field

result.tproperty.keyword 

Order by

Metric: Count 

Order Size


Descending  50



Group other values in separate bucket

Show missing values

Custom label


> Advanced

 Add

 Discard  Update


- i. **Aggregation:** select **Terms**.
- ii. **Field:** select **result.tclass.keyword**.
- iii. **Size:** set to **50**.
- iv. Click **Update** when finished.

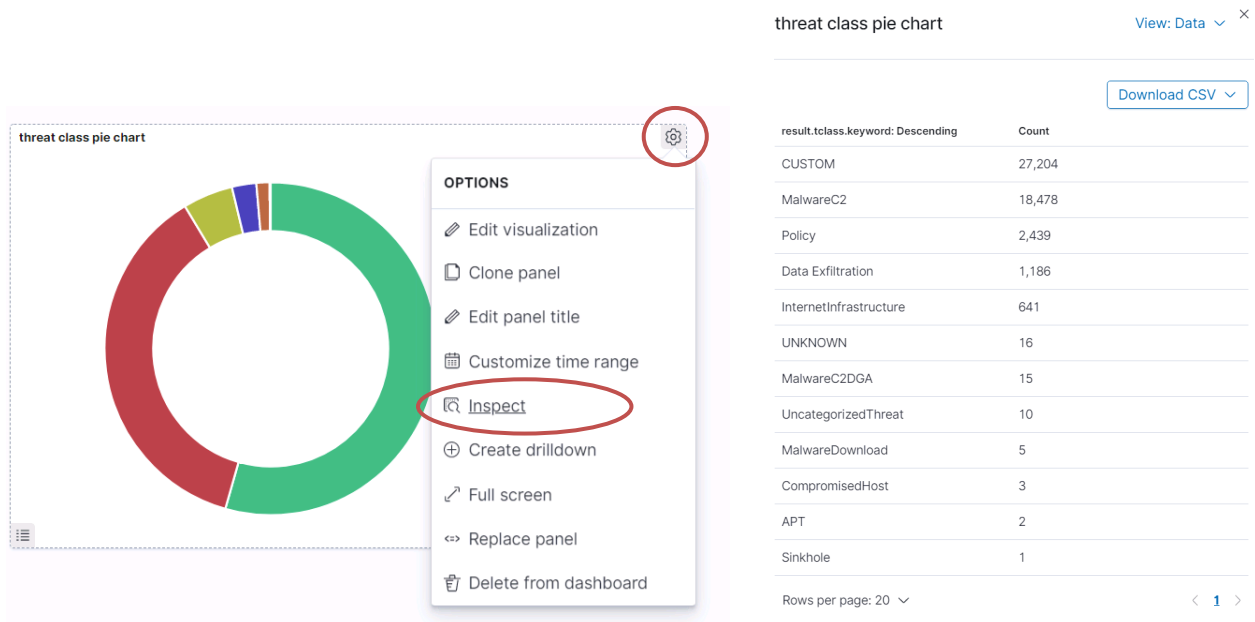
d. Your chart should look something like below. **Save** your chart.




e. Give it a **Title**. Click **Save and return**.

The screenshot shows a 'Save visualization' dialog box. The title of the dialog is 'Save visualization'. There is a 'Title' field with the text 'threat class pie chart'. Below the title field is a 'Description' field, which is currently empty. At the bottom of the dialog, there is a checkbox labeled 'Add to dashboards after saving' which is checked. Below the checkbox are two buttons: 'Cancel' and 'Save and return'.

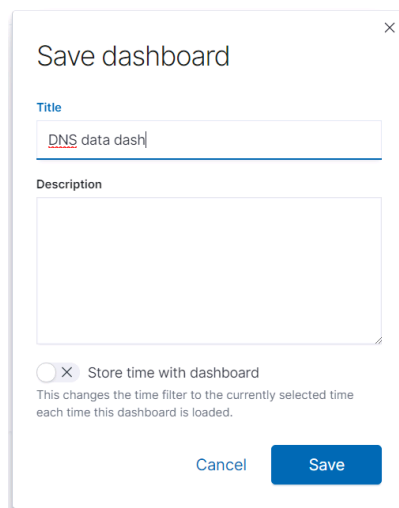
- f. Your chart now appears on your dashboard. Click the gears  options icon, then click **Inspect** to see a breakdown of data for the chart.



7. **Save** your dashboard.



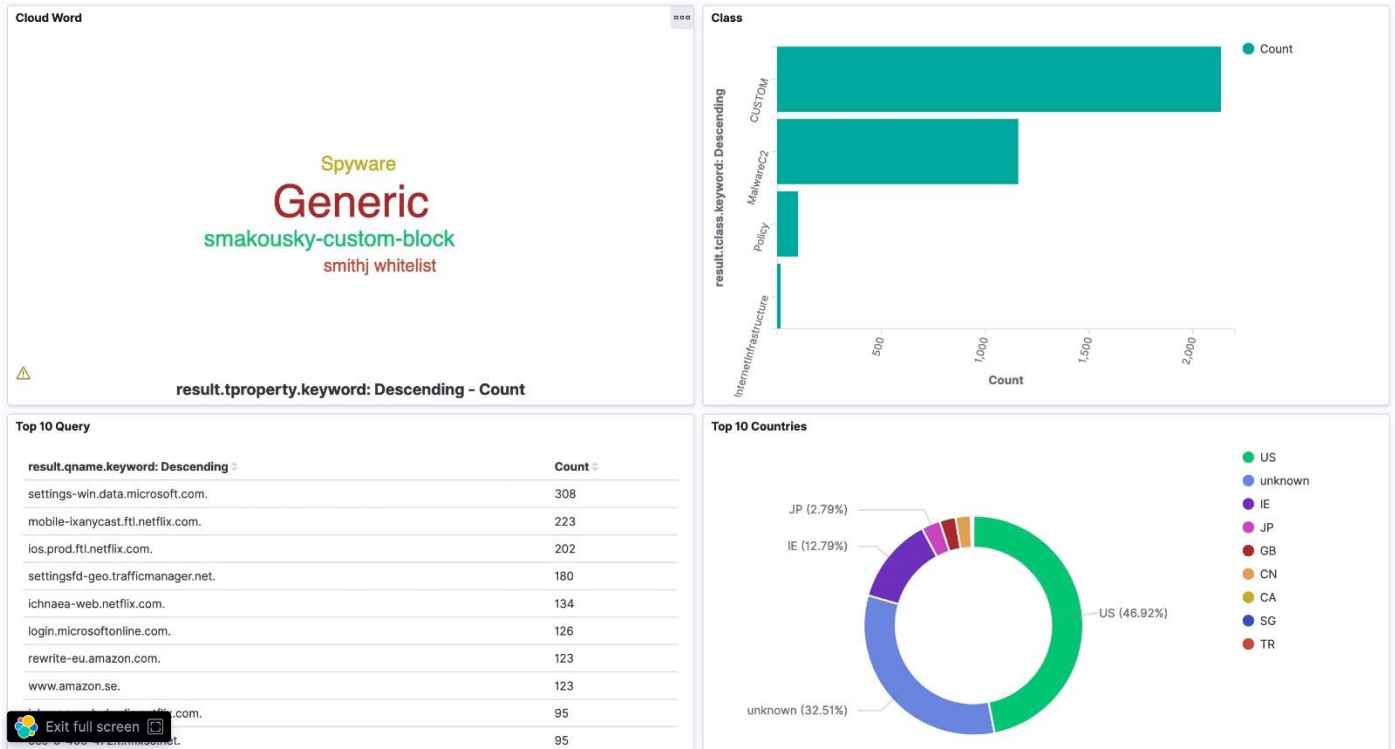
8. Give the dashboard a **Title**. Click **Save**.



The 'Save dashboard' dialog box contains the following fields and options:

- Title:** A text input field containing 'DNS data dash'.
- Description:** A larger text area for entering a description.
- Store time with dashboard**
This changes the time filter to the currently selected time each time this dashboard is loaded.
- Buttons: 'Cancel' and 'Save'.

You can add many different objects to dashboards, connect them, move them around, filter fields, and much more. Here are some examples showcasing more of Kibana's visualization capabilities:



Deployment Instructions: NIOS dnstap Logging

You can configure NIOS to use the dnstap log format to log DNS queries and responses at high rates to well-known destinations, such as an Ubuntu VM. dnstap is a flexible, structured binary log format for DNS software. It reduces the workload on NIOS to allow for logging queries and/or responses at higher speeds and performance than regular logging. This section shows you how to ingest dnstap formatted logs from NIOS into Elastic.

Several components are required to ingest dnstap logs into Elastic from NIOS. You need:

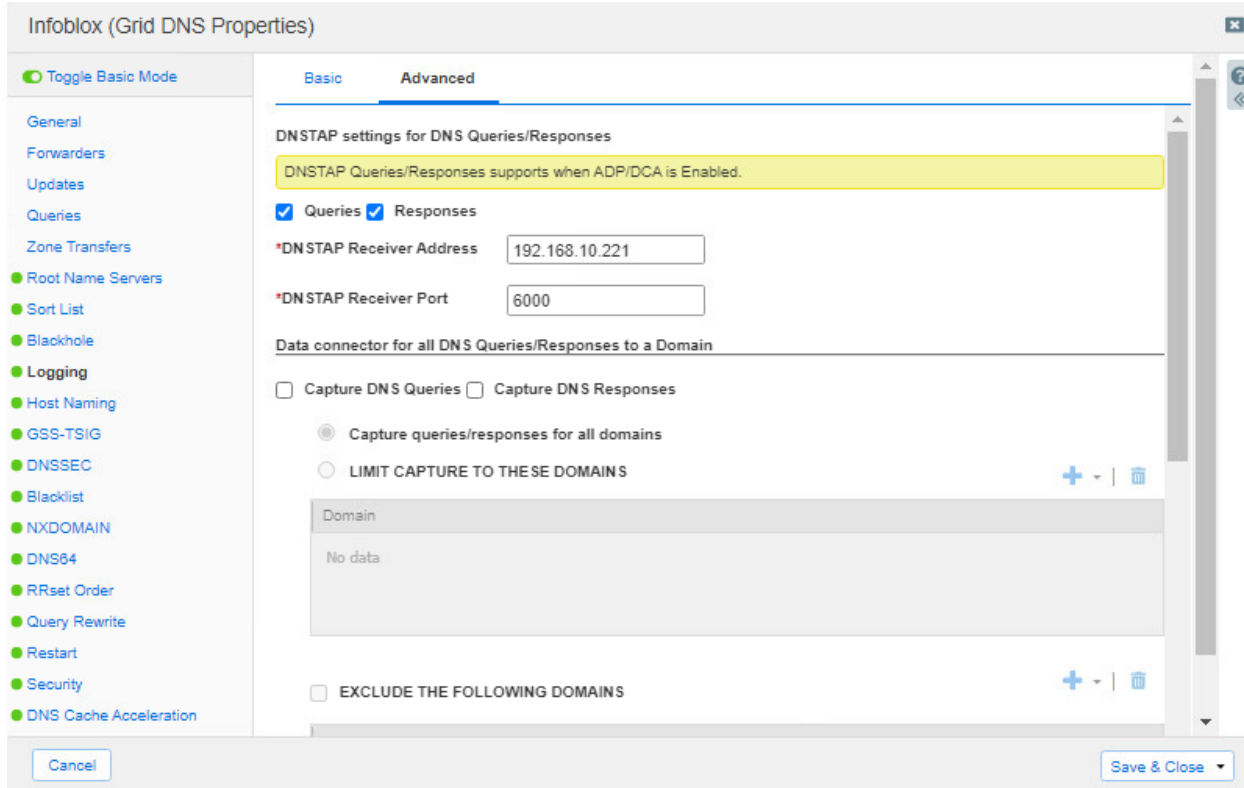
- NIOS with dnstap enabled. *For this demo, an IB-FLEX box was used with the DNS Cache Acceleration service running.*
- An external client to receive dnstap logs from NIOS. *For this demo, an Ubuntu 18.04 VM was used.*
- A way to receive, store and process dnstap logs after logging queries to the external client. *For this demo, the Python module [dnstap-receiver](#) was used.*
- Elastic Stack. *For this demo, Elastic Stack was installed on the same Ubuntu VM as dnstap-receiver.*

It is likely that if you are using dnstap, you are logging a large number of queries. It is highly recommended to secure DNS queries and responses between a server and client. Another layer of extra security that NIOS offers is the [DNS over TLS and DNS over HTTPS](#) services. These services encrypt DNS queries and responses to secure communication between a DNS server and a DNS client.

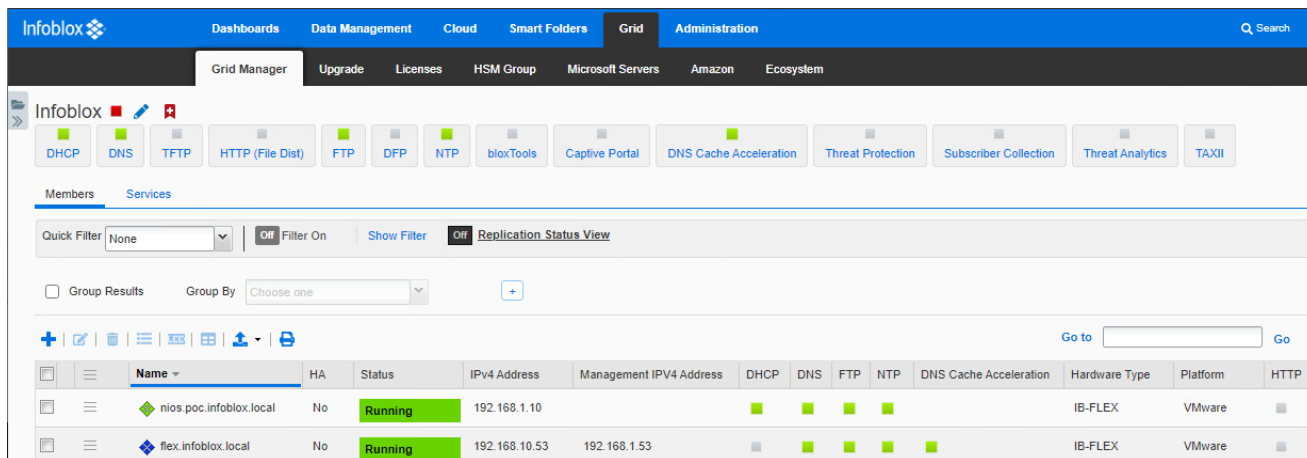
NIOS Configuration

dnstap is only available on NIOS 8.5.1 and higher and can only be configured on specific NIOS boxes running certain services. More information on requirements as well as detailed instructions on configuring dnstap for NIOS can be found on the official Infoblox documentation [here](#).

For this demo, dnstap was enabled on an IB-FLEX box running DNS Cache Acceleration (DCA). The following screenshot shows the NIOS configuration under the *Grid DNS Properties* editor. Observe that both queries and responses are being logged. The **DNSTAP Receiver Address** is the IP of the Ubuntu VM for which Elastic and the dnstap-receiver are installed. The default port used is **6000**.



For reference, the below screenshot shows the IP configuration and services enabled for the Grid.



dnstap-receiver Configuration

NIOS currently has no way to store or process dnstap logs after they leave the Grid. You will need some way to unpack and read the incoming dnstap messages from NIOS. A simple solution is to install [dnstap-receiver](#), a python module that receives dnstap messages and outputs them in a way Elastic can ingest. It supports several input stream types and can be configured to output readable data in many different ways, such as to a syslog server, stdout, a file, and more.

We will be using dnstap-receiver to ingest the dnstap messages from NIOS and output them to a file. Later, we will configure Logstash to ingest the file into Kibana. dnstap-receiver is configured with external config files,

similar to Logstash. Let's create the config file that will output the readable messages to a file. *Note: For this demo, dnstap-receiver was installed on the same Ubuntu VM as Elastic. It is best practice to keep these pieces of software installed on separate VMs.*

1. Access the machine where your dnstap-receiver instance is installed. Follow the installation instructions on its Python module page to install it.
2. Open a terminal.
3. To keep tidy, create a new directory for which dnstap will output the logfile that will be ingested by Elastic:

```
sudo mkdir /var/log/dnstap
```

4. Then create the logfile. *Note: We do this because the logfile must exist before executing dnstap-receiver. Otherwise it will throw an error.*

```
sudo touch /var/log/dnstap/dnstap.json
```

5. You must allow the logfile to be written to by dnstap-receiver. Enter the following command to allow all the files inside `/var/log/dnstap` to be readable, writable, and executable to all users on the computer. You can store the logfile in a writeable directory somewhere else, such as Documents, if you do not wish to change permissions.

```
sudo chmod -R 777 /var/log/dnstap/
```

6. Now create a new directory for which the config file will live:

```
sudo mkdir /etc/dnstap-receiver
```

7. Then create the config file:

```
sudo touch /etc/dnstap-receiver/dnstap.conf
```

8. Open the file with `gedit` for editing:

```
sudo gedit /etc/dnstap-receiver/dnstap.conf
```


9. Copy and paste the following into the file. Save and close the file when finished.

```
output:
  file:
    # enable or disable
    enable: true
    # format available text|json|yaml
    format: json
    # log file path or null to print to stdout
    file: /var/log/dnstap/dnstap.json
    # max size for log file
    file-max-size: 100M
    # number of max log files
    file-count: 10
```

This file tells dnstap-receiver to output the dnstap messages in json format to the `dnstap.json` file we created earlier. You can set various other parameters here, such as the max file size of the logfile or the max number of files to keep. These files can potentially become very large so adjust according to your needs.

Logstash Configuration for dnstap

Now let's configure Logstash to grab the data in `dnstap.json` logged by dnstap-receiver. View [the Logstash Configuration for TD](#) section of this document for more details on Logstash and config files.

1. Access the machine where your Logstash instance is installed. *Note: For this demo Elastic Stack was installed on Ubuntu 18.04.*
2. Open a terminal.
3. Navigate to where your Logstash configuration (`.conf`) files are located. In this demonstrative environment, these files are located in `/etc/logstash/conf.d`. Input the following command to navigate to the correct directory:

```
cd /etc/logstash/conf.d
```

4. Create a new file called `dnstap-nios.conf`:

```
sudo touch dnstap-nios.conf
```

5. Open the file with `gedit` for editing:

```
sudo gedit dnstap-nios.conf
```

6. Copy and paste the following into the file. Save and close the file when finished.

```
input {
  file {
    path => "/var/log/dnstap/*"
    codec => "json"
    mode => "tail"
    sinedb_path => "/dev/null"
  }
}
output {
  elasticsearch {
    hosts => ["localhost:9200"]
    index => "dnstap-nios"
  }
}
```

Note we are grabbing everything in the `/var/log/dnstap` directory we created earlier. Because `dnstap-receiver` is set to *append* the `dnstap.json` file with all the `dnstap` messages, we set the mode to `tail`.

7. Navigate to your home directory for Logstash. For this demo, this is `/usr/share/logstash/`. Input the following command to navigate to the correct directory:

```
cd /usr/share/logstash
```

8. Run Logstash with your new configuration:

```
sudo bin/logstash -f /etc/logstash/conf.d/dnstap-nios.conf
```

Allow several minutes of processing. The console will inform you if there are any syntax errors with your config file.

Alternatively, you can simply restart the Logstash service, but the console will not warn you of any errors with your config file:

```
sudo systemctl restart logstash
```

Run & Test Configuration

Let's run our new configuration.

First, we need to run dnstap-receiver with the config file created in the [dnstap-receiver Configuration](#) section of this document. Run this command on the Ubuntu VM during the entire time you wish to collect dnstap messages from NIOS.

1. Open a terminal.
2. Run dnstap-receiver using the config file as a parameter:

```
dnstap_receiver -c /etc/dnstap-receiver/dnstap.conf
```

The console will tell you if there are syntax errors with your config file.

Let's run a few test queries now. We will use the `dig` command to do so.

1. Open a new terminal or terminal tab without halting the terminal where dnstap-receiver is running.
2. Run a couple dig commands using the IP address of the NIOS Grid Member running dnstap. Try querying `infoblox.com`. *Note: For this demo, the IP used in the `dig` is the LAN Interface IP of the IB-FLEX box running DNS Cache Acceleration.*

```
dig @192.168.10.53 infoblox.com
```

```
infoblox@U-ElasticSearch:~$ dig @192.168.10.53 infoblox.com
; <<>> DiG 9.11.3-1ubuntu1.14-Ubuntu <<>> @192.168.10.53 infoblox.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 62529
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1220
; COOKIE: a88ae5a47219cd95d78b32c16035ef29680a98cd5c842439 (good)
;; QUESTION SECTION:
;infoblox.com.                IN      A

;; ANSWER SECTION:
infoblox.com.                30     IN      A      23.185.0.3

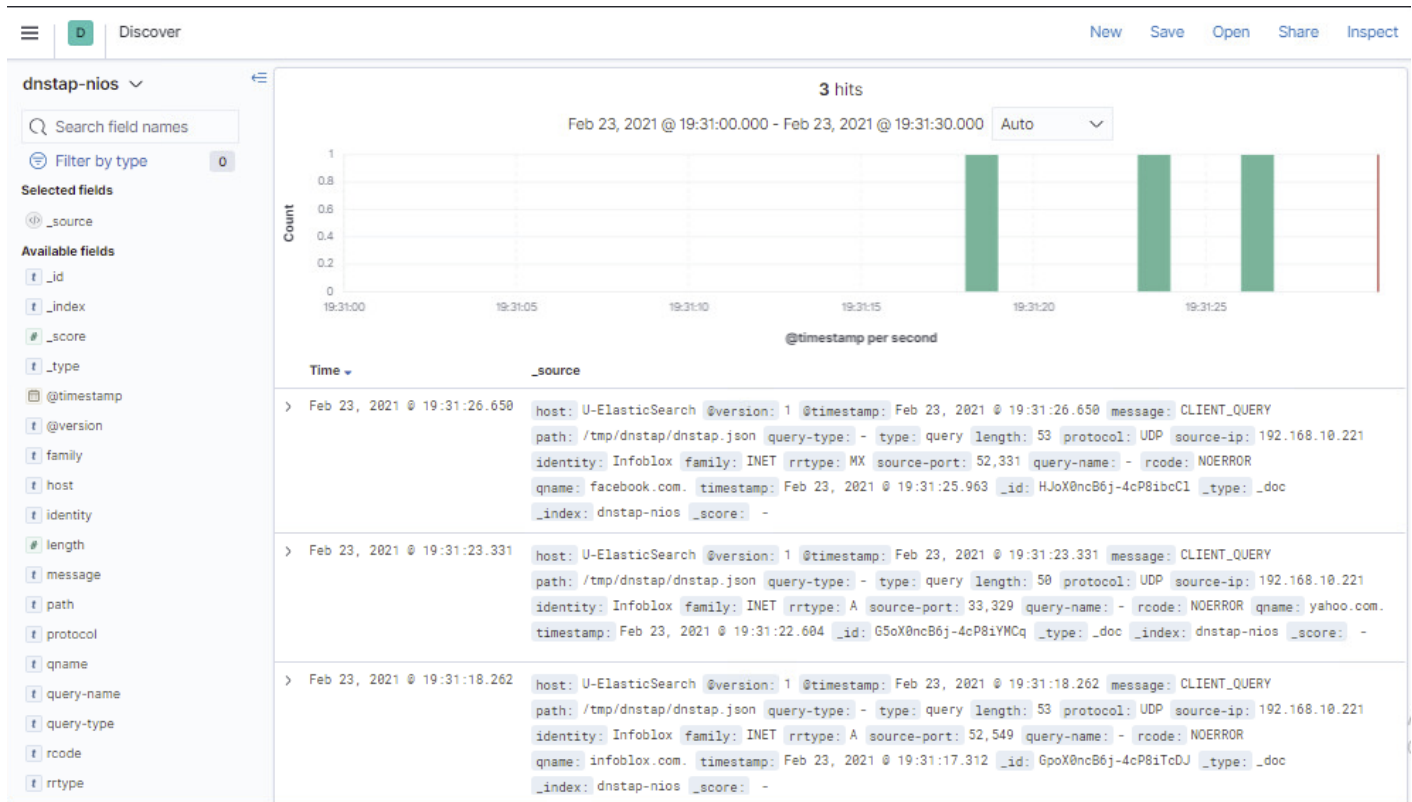
;; Query time: 12 msec
;; SERVER: 192.168.10.53#53(192.168.10.53)
;; WHEN: Tue Feb 23 22:16:09 PST 2021
;; MSG SIZE rcvd: 85

infoblox@U-ElasticSearch:~$
```

When queries are received by dnstap-receiver, they will appear in the dnstap-receiver stdout.

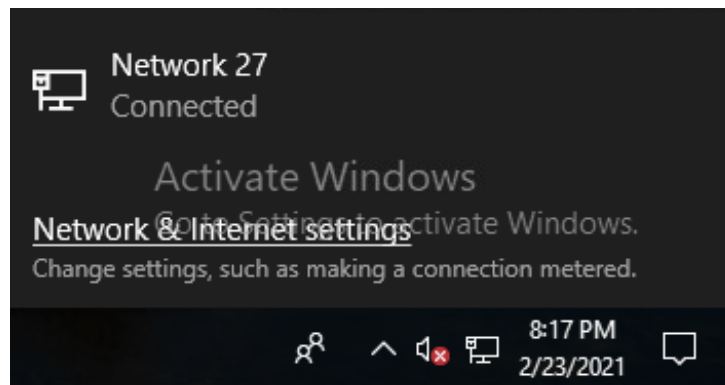
```
infoblox@U-ElasticSearch:~$ dnstap_receiver -c /etc/dnstap-receiver/dnstap.conf
2021-02-24T03:28:11.390426+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.10.221 36364 INET UDP 85b infoblox.com. A
2021-02-24T03:28:21.788973+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.10.221 57917 INET UDP 162b yahoo.com. A
2021-02-24T03:28:32.908814+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.10.221 35327 INET UDP 96b facebook.com. MX
```

- Open Kibana and observe the queries are ingested into Kibana. For more information on Kibana, indices and creating visualizations, see the [Kibana Data Discovery](#) and [Kibana Data Visualization](#) sections of this document.




We can also test with a Windows machine to see the ingesting in action instead of sending dig requests.

- Access a Windows machine on the same network as your Grid. *Note: For this demo, we will use the same Windows machine used to access the Kibana UI.*
- Open your **Internet and Network settings**.

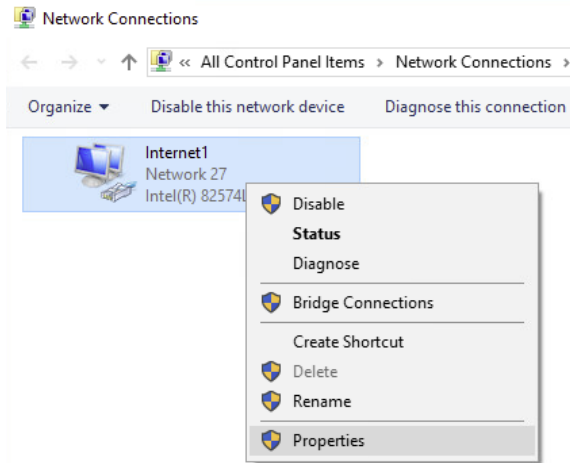


3. Click **Change adapter options**.

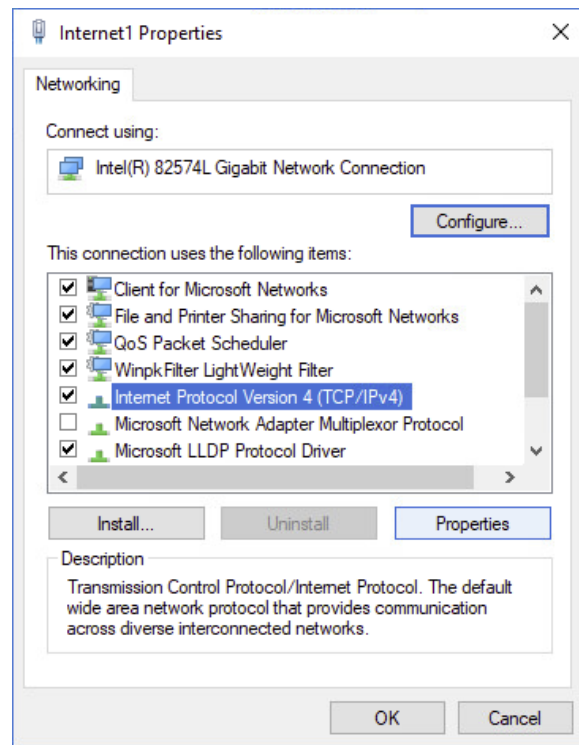
Change your network settings



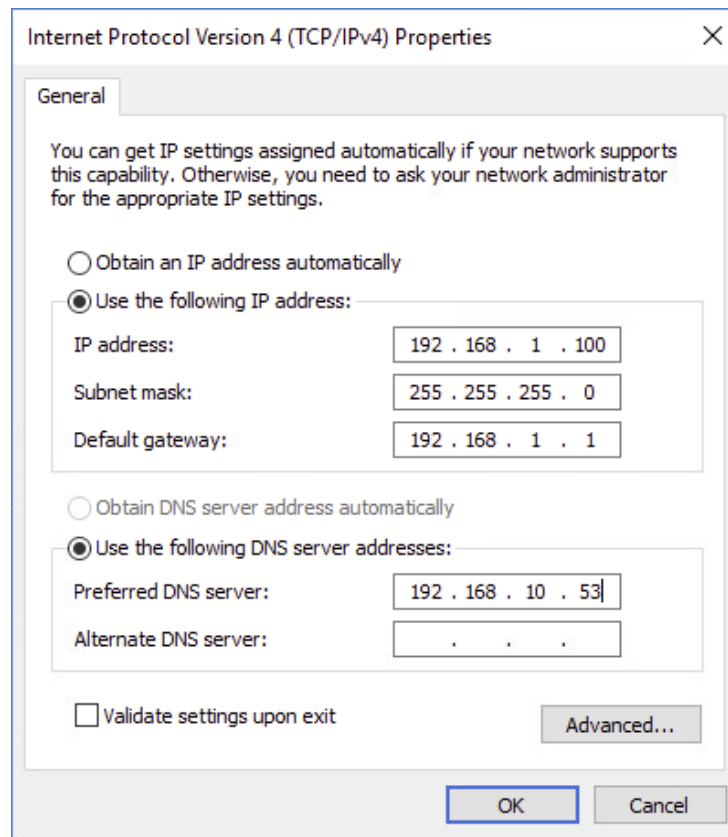
4. Right-click on your desired connection and click **Properties**.



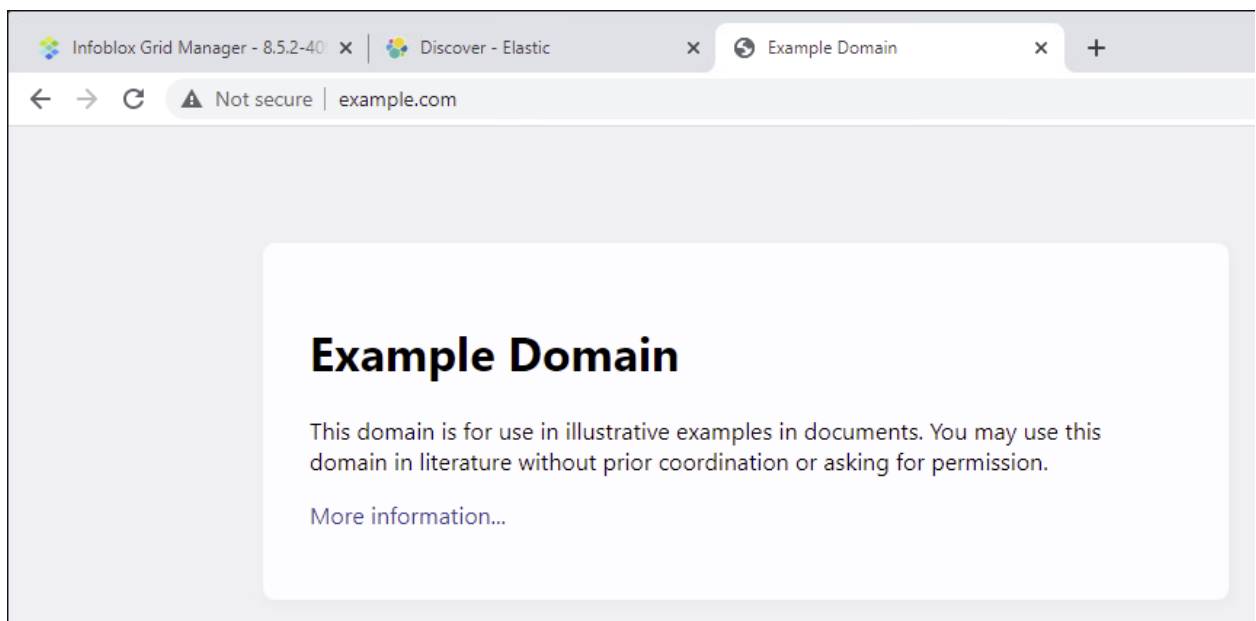
5. Click on **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.



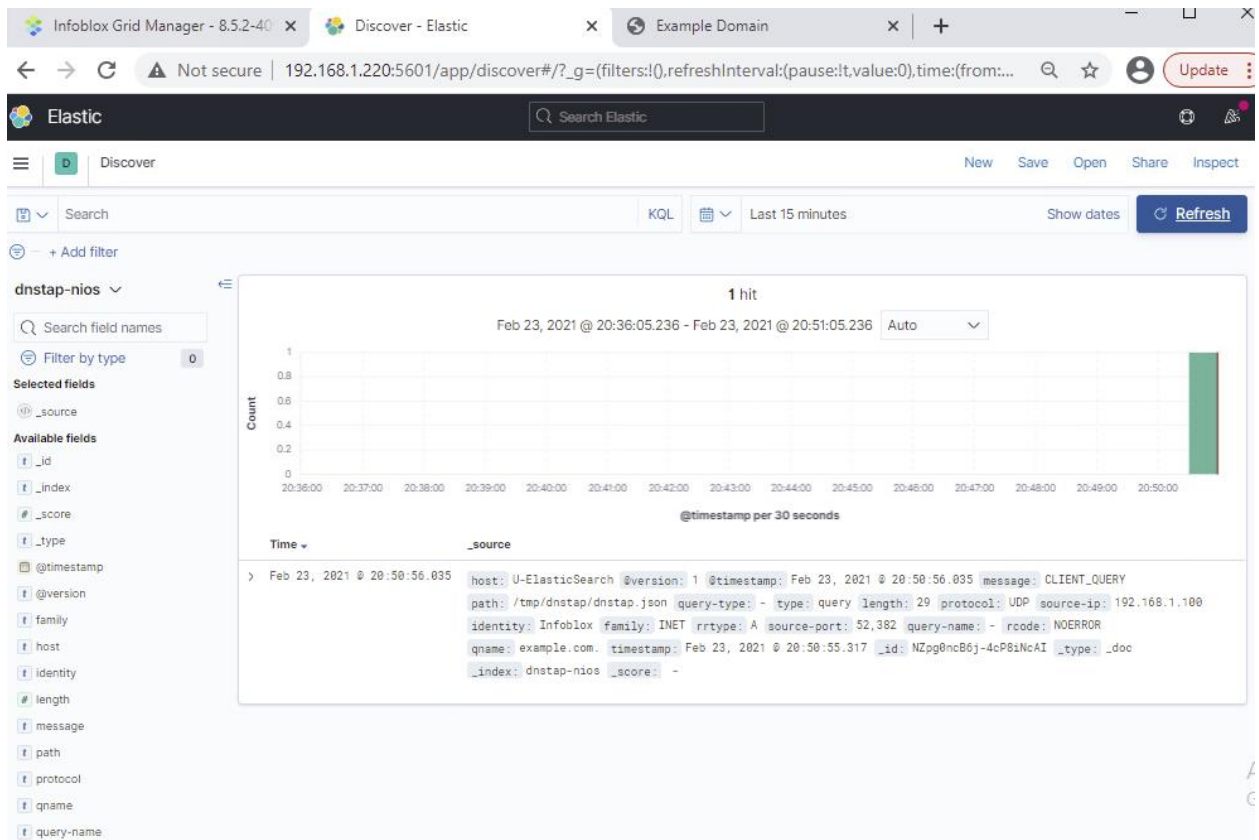
- Set the **Preferred DNS server** to the LAN IP address of the NIOS Grid Member running dnstap. Click **OK** when done. *Note: For this demo, the IP used is the LAN Interface IP of the IB-FLEX box running DNS Cache Acceleration.*



- Open Chrome. Go to <http://example.com>.



8. In Kibana, observe the query has been ingested.



The query also appears in the dnstap-receiver stdout.

```
^Cinfoblox@U-ElasticSearch:~$ dnstap_receiver -c /etc/dnstap-receiver/dnstap.conf
2021-02-24T05:02:15.208421+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.10.221 42915 INET UDP 85b infoblox.com. A
2021-02-24T05:02:21.111001+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.10.221 38949 INET UDP 162b yahoo.com. A
2021-02-24T05:02:24.309733+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.10.221 40881 INET UDP 96b facebook.com. MX
2021-02-24T05:03:49.563125+00:00 Infoblox CLIENT_RESPONSE NOERROR 192.168.1.100 51292 INET UDP 45b example.com. A
```



Infoblox enables next level network experiences with its Secure Cloud-Managed Network Services. As the pioneer in providing the world's most reliable, secure and automated networks, we are relentless in our pursuit of network simplicity. A recognized industry leader, Infoblox has 50 percent market share comprised of 8,000 customers, including 350 of the Fortune 500.

Corporate Headquarters | 3111 Coronado Dr. | Santa Clara, CA | 95054
+1.408.986.4000 | 1.866.463.6256 (toll-free, U.S. and Canada) | info@infoblox.com | www.infoblox.com



© 2020 Infoblox, Inc. All rights reserved. Infoblox logo, and other marks appearing herein are property of Infoblox, Inc. All other marks are the property of their respective owner(s).