Synapse Bootcamp - Module 9

Pivoting and Traversal in Storm - Exercises

Objectives

In these exercises you will learn:

- How to perform common pivot operations using Storm
- How to perform common traversal operations using Storm
- How the ability to examine large data sets differs when using UI navigation vs. Storm

> **Note:** We are constantly updating Synapse and its Power-Ups! We do our best to make sure our course documents (slides, exercises, and answer keys) are up-to-date. However, you may notice small differences (such as between a screen capture in the documents and the appearance of your current instance of Synapse).
>
> If something is unclear or if you identify an error, please reach out to us so we can assist!

Exercises

- All exercises use the **Research Tool** with the **Storm Mode Selector** set to **Storm mode.**
- Some example queries may wrap due to length.

The **Storm Quick Reference** on Pivots and Traversals (included with the supplemental materials provided for this course) may be helpful for this (and future) exercises.

The online **pivot and traversal** reference includes detailed documentation and examples for all pivot and traversal operations. It is part of the Storm Reference included with the Synapse User Guide.


Pivoting

Exercise 1

Objectives:
- **Write and execute basic Storm queries using pivot operations.**
- **Use the special "pivot to tags" operation.**
- **Leverage the 'uniq' Storm command to deduplicate results.**

Note: You can use either **explicit** or **implicit** pivot syntax for the following exercise.

You are researching a suspicious IPv4 address and want to know what FQDNs have resolved to that IP.

- In the **Research Tool,** ensure your **Storm Query Bar** is in **Storm mode** and your display mode is set to **Tabular:**
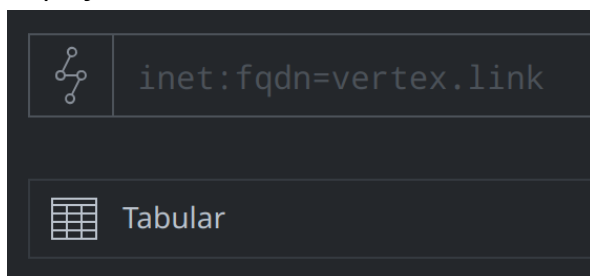
- Enter the following in the **Storm Query Bar** and press **Enter** to run the query:

  `inet:ipv4=173.254.222.138`

  Question 1: How can you **add** to the Storm query in order to **pivot** to the associated DNS A (`inet:dns:a`) records? How many DNS A records are returned?

  Question 2: How can you **add** to your Storm query to **pivot** from the DNS A records to the associated FQDNs (`inet:fqdn` nodes)? How many FQDNs are returned?

  Question 3: How can you **add** to your Storm query to **pivot to tags** and view the `syn:tag` nodes for the tags on the FQDNs?

  Question 4: How many `syn:tag` nodes are returned? Why are there duplicates?

  Question 5: What Storm command can you **add** to your Storm query to **deduplicate** ("unique") your results and only display one instance of each tag?

Exercise 2

Objectives:
- **Write and execute basic Storm queries using pivot operations.**
- **Compare and contrast use of the Synapse UI and use of Storm to examine data.**
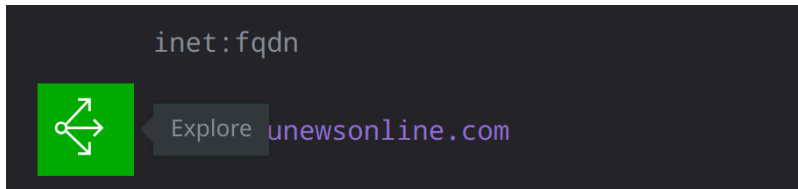
You are researching malicious FQDNs and want to identify any files (malware samples) that communicate with those FQDNs.

- Enter the following in the **Storm Query Bar** and press **Enter** to run the query:
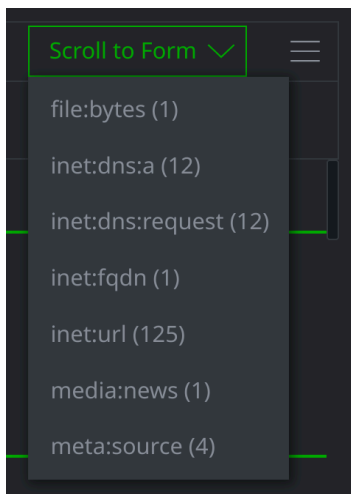
  `inet:fqdn=media.aunewsonline.com`

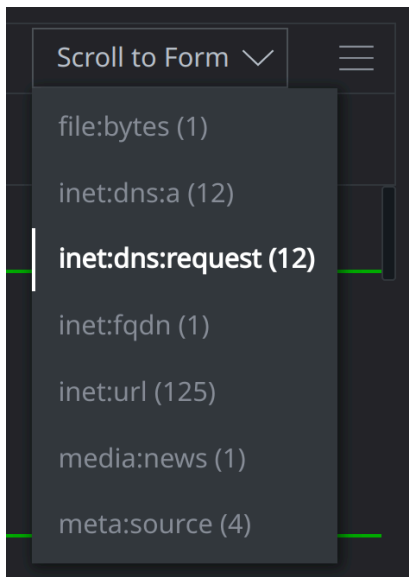First, we'll use the **Explore** button to answer this question.

- Click the **Explore** button next to the FQDN to display adjacent nodes:
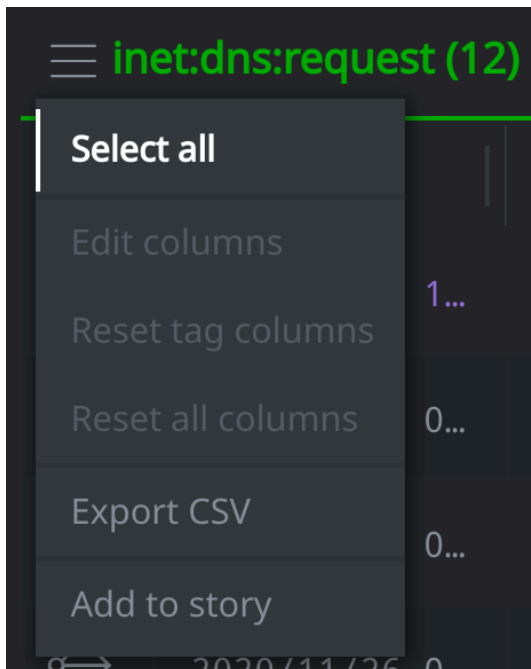

- Click the **Scroll to Form** button. **Note** the number and kinds of forms in the dropdown list:
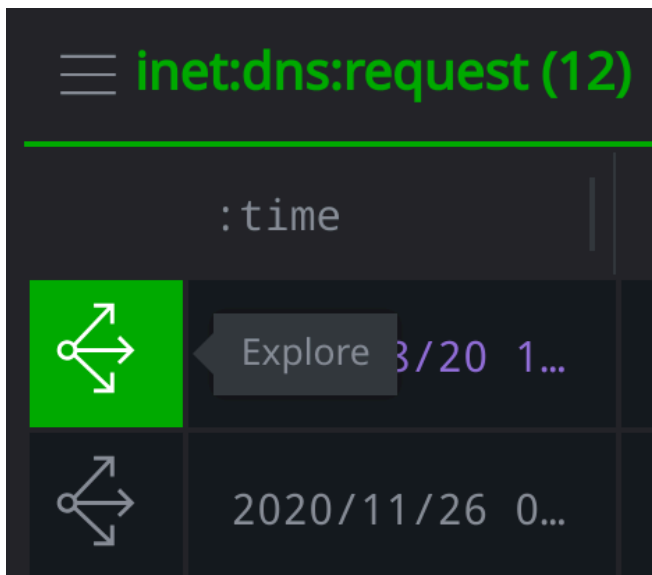

- Choose **inet:dns:request** from the dropdown list:

- Click the **hamburger menu** next to the **inet:dns:request** table header and choose **Select All:**

- Click the **Explore** button next to any of the `inet:dns:request` nodes to display adjacent nodes:

Question 1: How many files perform DNS queries for the FQDN `media.aunewsonline.com`?

It was easy to answer the question with the UI! But you had to navigate (Explore) through a lot of nodes that were not related to your question.

Using Storm to **pivot** directly to the data you need can be more efficient when you are asking a specific question. Let's answer the same question using Storm.

Note: You can use either **explicit** or **implicit** pivot syntax for the following exercise.

- Click **query** in your **breadcrumbs** to return to your original query:
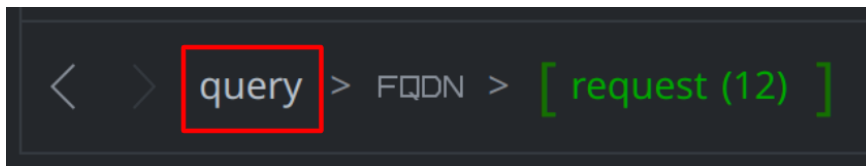

```
inet:fqdn=media.aunewsonline.com
```

Question 2: How can you **add** to your Storm query to **pivot** to the associated DNS requests (`inet:dns:request` nodes)?

Question 3: How can you **add** to your Storm query to **pivot** from the DNS requests to the files (`file:bytes`) that make the requests?

Question 4: Do you have duplicate results? If so, how can you **add** to your query to remove deduplicate ("unique") the results?

Edge Traversal

Exercise 3

Objective:
- **Write and execute basic Storm queries using edge traversal operations.**

You are researching a malicious FQDN. You want to know if it has appeared in any articles (`media:news` nodes) that might provide more context.

- Enter the following into your **Storm Query Bar** and press **Enter** to run the query:

```
inet:fqdn=chemscalere.com
```

Question 1: How can you **add** to your query to **traverse any light edges** and find **any** nodes that "point to" the FQDN?

> **Hint:** Recall that all lightweight edges have a **direction.** If we want things that "point to" the FQDN, the traversal "arrow" should face to the **left.**

Question 2: What kinds of nodes are linked to the FQDN using light edges?

> We want to see articles (`media:news` nodes) that include our FQDN. Our results include other nodes that we are not interested in.

Question 3: How can you **modify** your Storm query to return **only** articles (`media:news` nodes) that include the FQDN?

Using Synapse with Large Data Sets

Exercise 4

> **Objectives:**
> - **Write and execute basic Storm queries using pivot operations for large data sets.**
> - **Compare and contrast using the Synapse UI and using Storm to examine data.**

In **Exercise 2**, we used both the Synapse UI and Storm to find the files that query the malicious FQDN `media.aunewsonline.com`.

In that example, there was not much difference between the two methods.

However, there are times when using **Explore** is less efficient than using Storm. Sometimes this difference is significant! This is especially true when performing queries using **large data sets.**

In Exercise 2 we started with one FQDN (`media.aunewsonline.com`) and found malware that queries that FQDN.

What if, instead of asking about **one** FQDN, we want to find malware that queries **any** APT1 FQDN? (There are just over 2,000 FQDNs.)

First let's try to answer the question using the **Explore** button.

- Enter the following in the **Storm Query Bar** and press **Enter** to run the query to lift **all** of the APT1 FQDNs:

```
inet:fqdn#rep.mandiant.apt1
```

- Click the **hamburger menu** to the left of the **inet:fqdn** table header and choose **Select All** to select all 2,073 results:
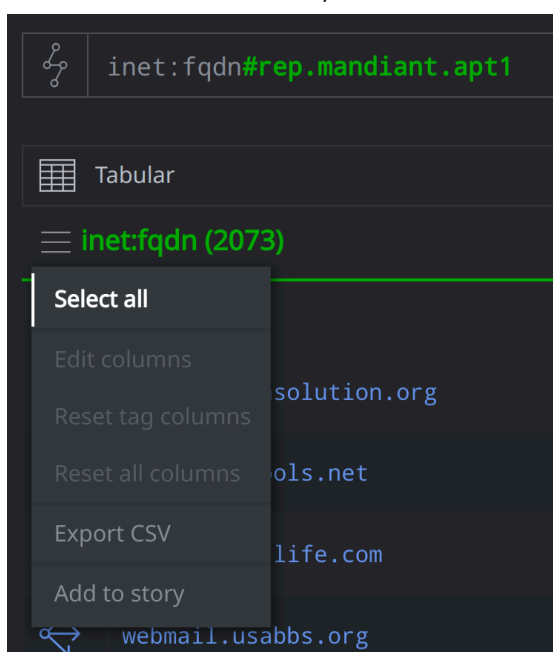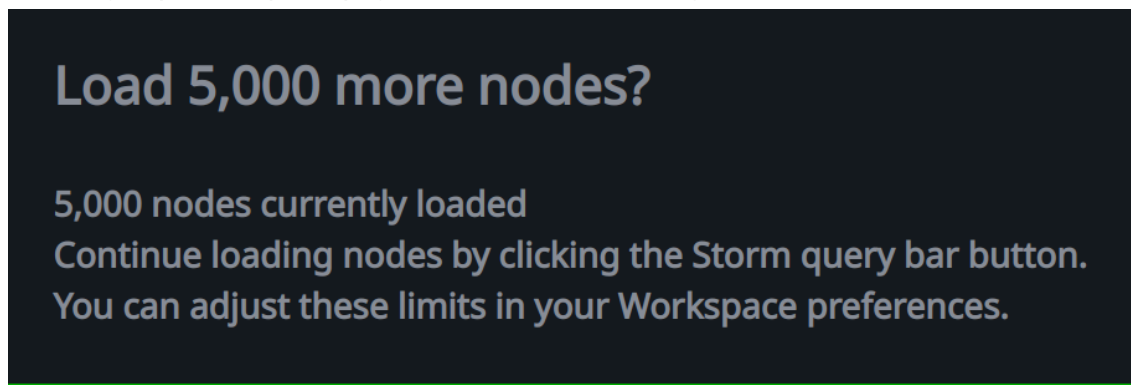
- Click the **Explore** button next to any of the `inet:fqdn` nodes to display adjacent nodes:
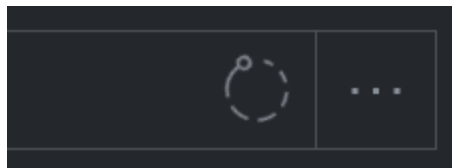

Note: Synapse will prompt you to continue loading results:

You can:
- choose **not** to continue loading, or
- **cancel** a running query at any time by clicking the spinning **Query Status Icon** to the right of the **Storm Query Bar:**

Question 1: What happens? Are you able to use the **Explore** button, or do you get stuck? (That is, do you get tired of loading results and wondering when Synapse will finish?)

Now we'll answer the question using a **Storm query.**

- Enter the following into your **Storm Query Bar** and press **Enter** to run the query:

```
inet:fqdn#rep.mandiant.apt1 -> inet:dns:request -> file:bytes
   | uniq
```

Question 2: What happens? How many files are returned?