

Synapse Bootcamp - Module 9

Pivoting and Traversal in Storm - Answer Key

Pivoting and Traversal in Storm - Answer Key	1
Answer Key	2
Pivoting	2
Exercise 1 Answer	2
Exercise 2 Answer	5
Edge Traversal	9
Exercise 3 Answer	9
Using Synapse with Large Data Sets	11
Exercise 4 Answer	11

Answer Key

Pivoting

Exercise 1 Answer

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.

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?

- Add a **pivot** operation to your existing query:

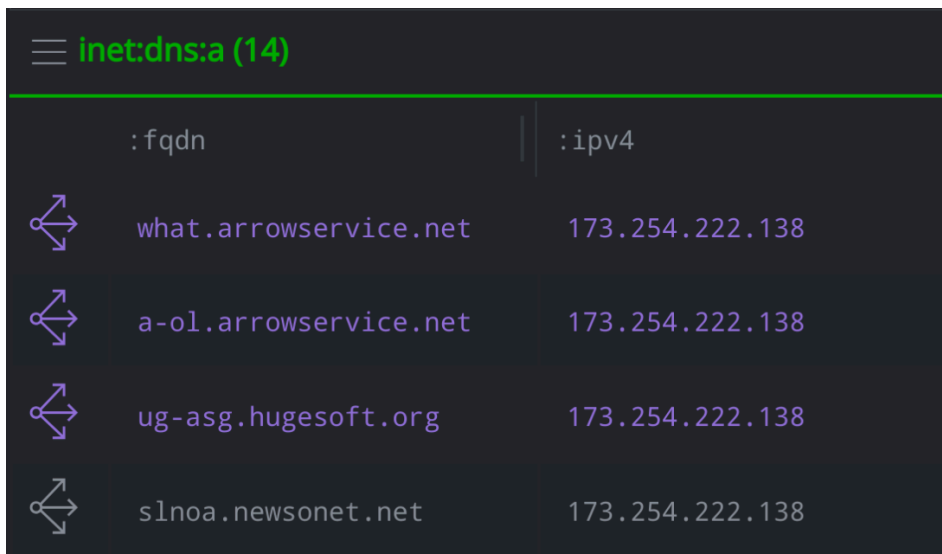
Explicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a:ipv4
```

Implicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a
```

- There are **14** **inet:dns:a** nodes returned by the query:



inet:dns:a (14)	
:fqdn	:ipv4
what.arrowservice.net	173.254.222.138
a-01.arrowservice.net	173.254.222.138
ug-asg.hugesoft.org	173.254.222.138
s1noa.newsonet.net	173.254.222.138

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?

- Add a **pivot** operation to your existing query:

Explicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a:ipv4 :fqdn -> inet:fqdn
```

Implicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a -> inet:fqdn
```

- There are **14** **inet:fqdn** nodes returned by the query:



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?

- Add another **pivot** operation to your existing query. In order to **pivot to tags**, use the special hashtag (**#**) character as the target of the pivot:

Explicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a:ipv4 :fqdn -> inet:fqdn  
-> #
```

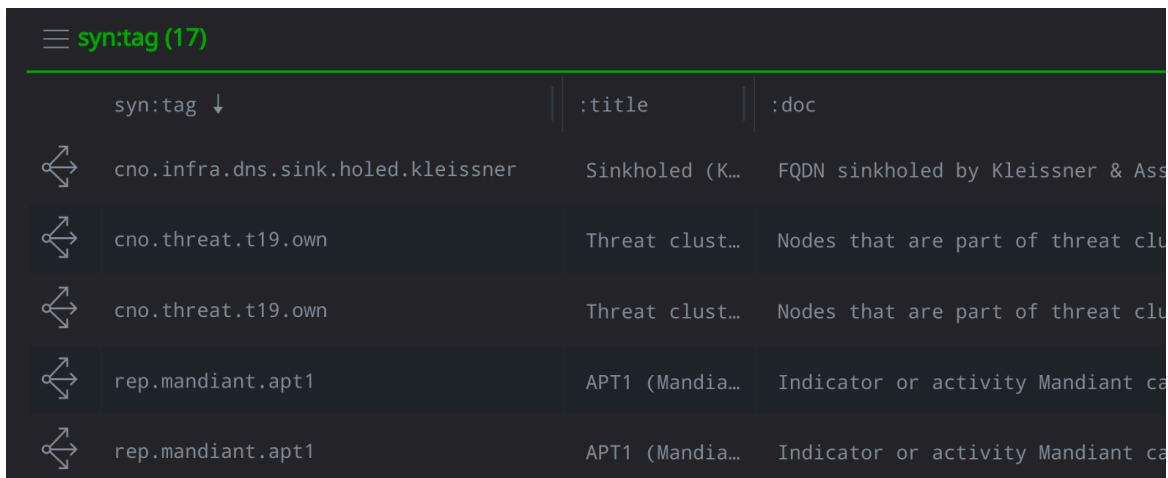
Implicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a -> inet:fqdn -> #
```

Because "pivot to tags" is a special pivot operation (it pivots specifically from **tags** on nodes to **syn:tag** nodes), there is no "explicit" vs "implicit" syntax for this operation; it is the same in both cases.

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

- The query returns **17 syn:tag** nodes:



syn:tag (17)		
syn:tag ↓	:title	:doc
cno.infra.dns.sink.holed.kleissner	Sinkholed (K...	FQDN sinkholed by Kleissner & Ass
cno.threat.t19.own	Threat clust...	Nodes that are part of threat clu
cno.threat.t19.own	Threat clust...	Nodes that are part of threat clu
rep.mandiant.apt1	APT1 (Mandia...	Indicator or activity Mandiant ca
rep.mandiant.apt1	APT1 (Mandia...	Indicator or activity Mandiant ca

There are **duplicate** results because many of the **inet:fqdn** nodes (the **source** nodes for your pivot to tags operation) have the **same tag** applied to each node.

For example: all 14 FQDNs have the tag **rep.mandiant.apt1**. When you pivot in Storm, Synapse returns one instance of the **syn:tag** node for **each** FQDN that has the tag. So there are 14 "copies" of the **rep.mandiant.apt1** tag in your results.

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?

- Add the **uniq** command to the end of your query to "unique" the results:

Explicit syntax:

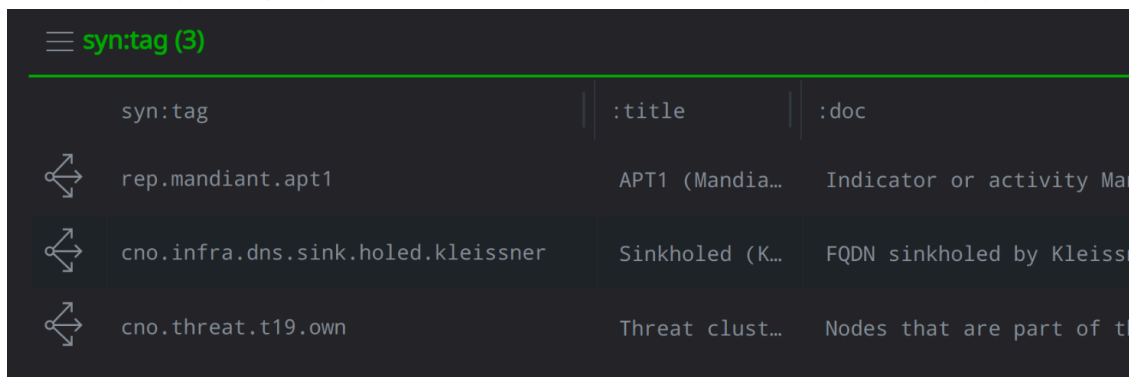
```
inet:ipv4=173.254.222.138 -> inet:dns:a:ipv4 :fqdn -> inet:fqdn  
-> # | uniq
```

Implicit syntax:

```
inet:ipv4=173.254.222.138 -> inet:dns:a -> inet:fqdn -> #  
| uniq
```

Recall that you use the pipe character (|) to switch from a Storm query to a Storm command and vice versa.

After running this query you should have only **three syn:tag** nodes in your results:



The screenshot shows a table of search results for the query 'syn:tag (3)'. The table has three columns: 'syn:tag', ':title', and ':doc'. There are three rows of results, each with a left-pointing arrow icon next to the 'syn:tag' column.

syn:tag	:title	:doc
rep.mandiant.apt1	APT1 (Mandia...	Indicator or activity Ma
cno.infra.dns.sink.holed.kleissner	Sinkholed (K...	FQDN sinkholed by Kleiss
cno.threat.t19.own	Threat clust...	Nodes that are part of t

Running the **uniq** command is not affected by using either explicit or implicit syntax; the command can be added to either type of query.

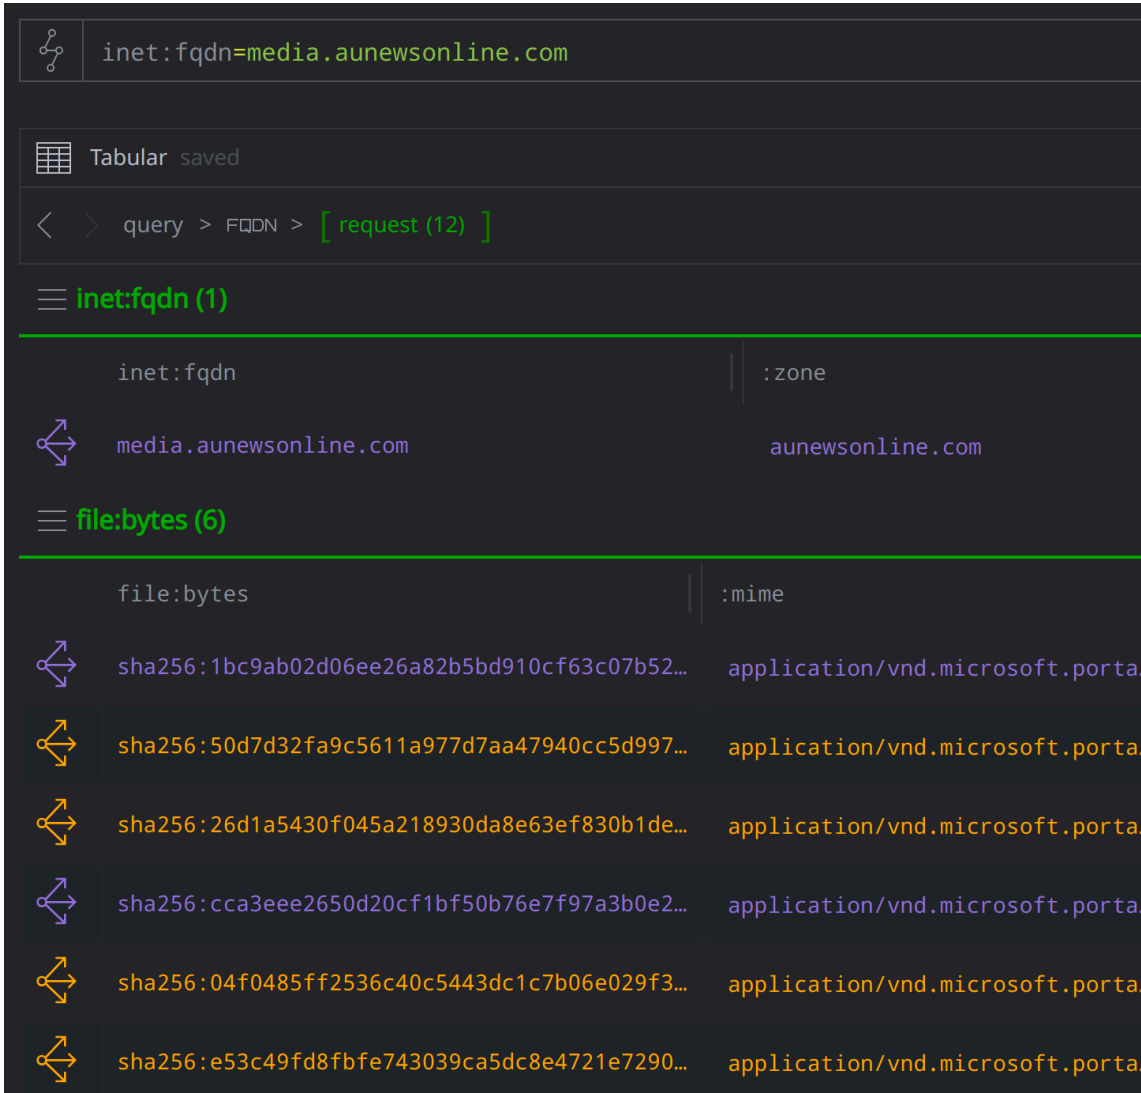
Exercise 2 Answer

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.

Question 1: How many files perform DNS queries for the FQDN [media.aunewsonline.com](https://www.aunewsonline.com)?

- There are **six** files that make DNS queries for that FQDN:



The screenshot shows a Storm query interface with the following structure:

- Query: `inet:fqdn=media.aunewsonline.com`
- Navigation: `query > FQDN > [request (12)]`
- Section: `inet:fqdn (1)`
- Table 1:

inet:fqdn	:zone
media.aunewsonline.com	aunewsonline.com
- Section: `file:bytes (6)`
- Table 2:

file:bytes	:mime
sha256:1bc9ab02d06ee26a82b5bd910cf63c07b52...	application/vnd.microsoft.porta...
sha256:50d7d32fa9c5611a977d7aa47940cc5d997...	application/vnd.microsoft.porta...
sha256:26d1a5430f045a218930da8e63ef830b1de...	application/vnd.microsoft.porta...
sha256:cca3eee2650d20cf1bf50b76e7f97a3b0e2...	application/vnd.microsoft.porta...
sha256:04f0485ff2536c40c5443dc1c7b06e029f3...	application/vnd.microsoft.porta...
sha256:e53c49fd8fbfe743039ca5dc8e4721e7290...	application/vnd.microsoft.porta...

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

- Add a **pivot** operation to your query:

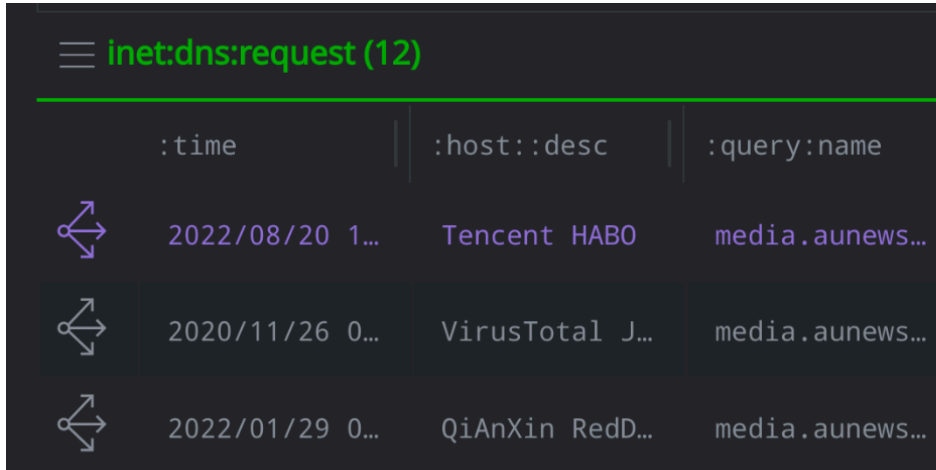
Explicit syntax:




```
inet:fqdn=media.aunewsonline.com
-> inet:dns:request:query:name:fqdn
```

Implicit syntax:

```
inet:fqdn=media.aunewsonline.com -> inet:dns:request
```

You should see **12 inet:dns:request** nodes:



	:time	:host::desc	:query:name
	2022/08/20 1...	Tencent HABO	media.aunews...
	2020/11/26 0...	VirusTotal J...	media.aunews...
	2022/01/29 0...	QiAnXin RedD...	media.aunews...

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?

- Add another **pivot** operation to your query:

Explicit syntax:

```
inet:fqdn=media.aunewsonline.com  
-> inet:dns:request:query:name:fqdn :exe -> file:bytes
```

Implicit syntax:

```
inet:fqdn=media.aunewsonline.com -> inet:dns:request  
-> file:bytes
```

You should see **24 file:bytes** nodes:

```
≡ file:bytes (24)
file:bytes | :mime
↔ sha256:1bc9ab02d06ee26a82b5bd910cf63c07b52... application/vnd.microsoft
↔ sha256:1bc9ab02d06ee26a82b5bd910cf63c07b52... application/vnd.microsoft
↔ sha256:50d7d32fa9c5611a977d7aa47940cc5d997... application/vnd.microsoft
```

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

- Add the **uniq** command to the end of your query to "unique" the results:

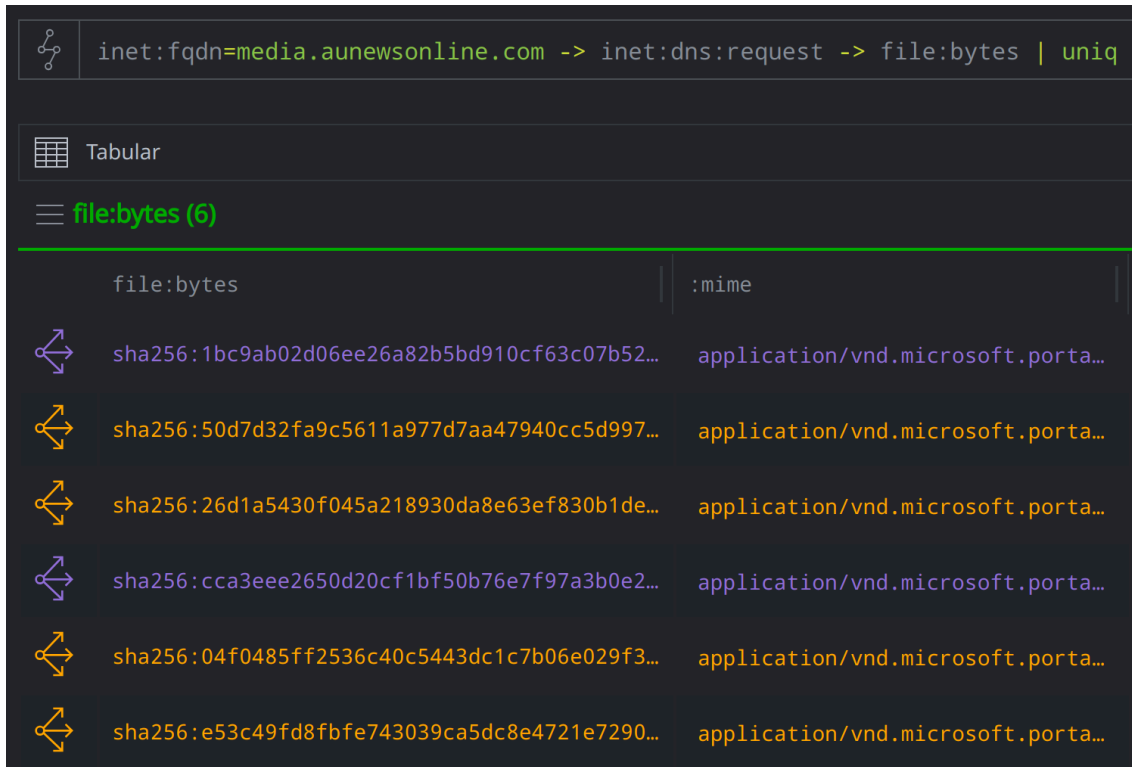
Explicit syntax:

```
inet:fqdn=media.aunewsonline.com
-> inet:dns:request:query:name:fqdn :exe -> file:bytes | uniq
```

Implicit syntax:

```
inet:fqdn=media.aunewsonline.com -> inet:dns:request
-> file:bytes | uniq
```


You should have **six** files - the same answer you received using the **Explore** button:



The screenshot shows a Synapse query interface. At the top, a query is entered: `inet:fqdn=media.aunewsonline.com -> inet:dns:request -> file:bytes | uniq`. Below the query, the view is set to "Tabular". The results are displayed as a table with 6 rows. The first row shows the column headers: "file:bytes" and ":mime". The subsequent rows show SHA256 hashes and their corresponding MIME types, all of which are "application/vnd.microsoft.porta...".

file:bytes	:mime
sha256:1bc9ab02d06ee26a82b5bd910cf63c07b52...	application/vnd.microsoft.porta...
sha256:50d7d32fa9c5611a977d7aa47940cc5d997...	application/vnd.microsoft.porta...
sha256:26d1a5430f045a218930da8e63ef830b1de...	application/vnd.microsoft.porta...
sha256:cca3eee2650d20cf1bf50b76e7f97a3b0e2...	application/vnd.microsoft.porta...
sha256:04f0485ff2536c40c5443dc1c7b06e029f3...	application/vnd.microsoft.porta...
sha256:e53c49fd8fbfe743039ca5dc8e4721e7290...	application/vnd.microsoft.porta...

Edge Traversal

Exercise 3 Answer

Objective:

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

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

- To find any nodes that "point to" our FQDN using a light edge, we can use a **wildcard edge traversal** operation:

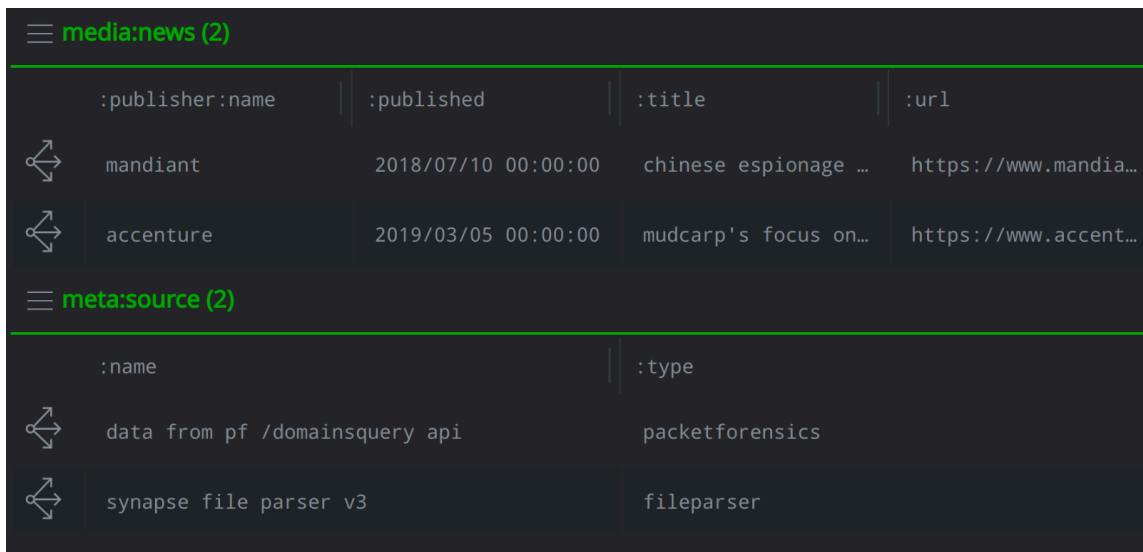
```
inet:fqdn=chemscalere.com <(*)- *
```



Note that this syntax uses two wildcard characters (*):



- One in the edge traversal operator to indicate Synapse should traverse **any** light edge (<(*)-).
- One representing the **target** forms for the operation (i.e., any / all forms).

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

- The FQDN is linked to:
 - Articles (**media:news** nodes)
 - Data sources (**meta:source** nodes)



media:news (2)				
	:publisher:name	:published	:title	:url
	mandiant	2018/07/10 00:00:00	chinese espionage ...	https://www.mandia...
	accenture	2019/03/05 00:00:00	mudcarp's focus on...	https://www.accent...

meta:source (2)		
	:name	:type
	data from pf /domainsquery api	packetforensics
	synapse file parser v3	fileparser

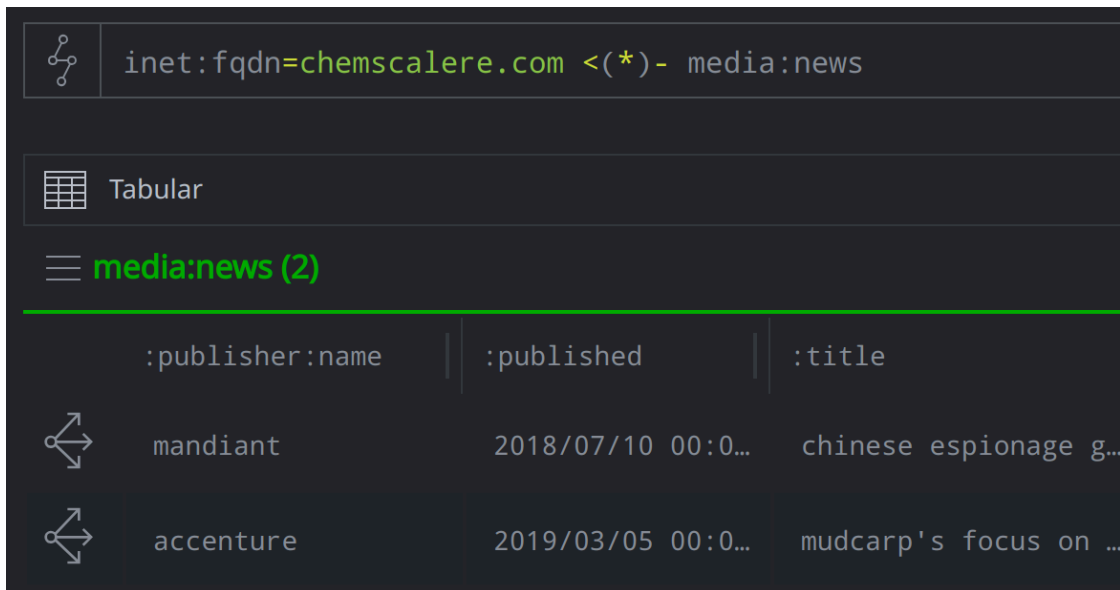
Note: This answer is based on the baseline Synapse demo instance. Your answer may vary depending on any data that has been added to your demo instance so far in this course.

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



- To return only the **media:news** nodes, we can change our target from a wildcard to only the nodes we want to see:

```
inet:fqdn=chemscalere.com <(*)- media:news
```

We are still traversing "any" edge using the edge wildcard (<(*)-). But we are telling Synapse that the "target" nodes are articles (**media:news** nodes):



The screenshot shows a Synapse interface with a Storm query: `inet:fqdn=chemscalere.com <(*)- media:news`. Below the query, the view is set to 'Tabular' and shows 2 results for the `media:news` node.

	:publisher:name	:published	:title
	mandiant	2018/07/10 00:0...	chinese espionage g...
	accenture	2019/03/05 00:0...	mudcarp's focus on ...

The following queries will **also** work:

- Tell Synapse you want to see **any** target nodes connected by a **refs** edge:

```
inet:fqdn=chemscalere.com <(refs)- *
```

meta:source nodes are connected by a **seen** edge, so the above query will only show us the **media:news** nodes.

- Tell Synapse you want to see the **media:news** nodes connected by **refs** edges:

```
inet:fqdn=chemscalere.com <(refs)- media:news
```

Using Synapse with Large Data Sets

Exercise 4 Answer

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.

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?)

- You get "stuck".

Synapse will load results as long as you allow it to continue. But your browser will become less and less responsive as you load more and more results. Synapse can handle large data sets - but your browser cannot.

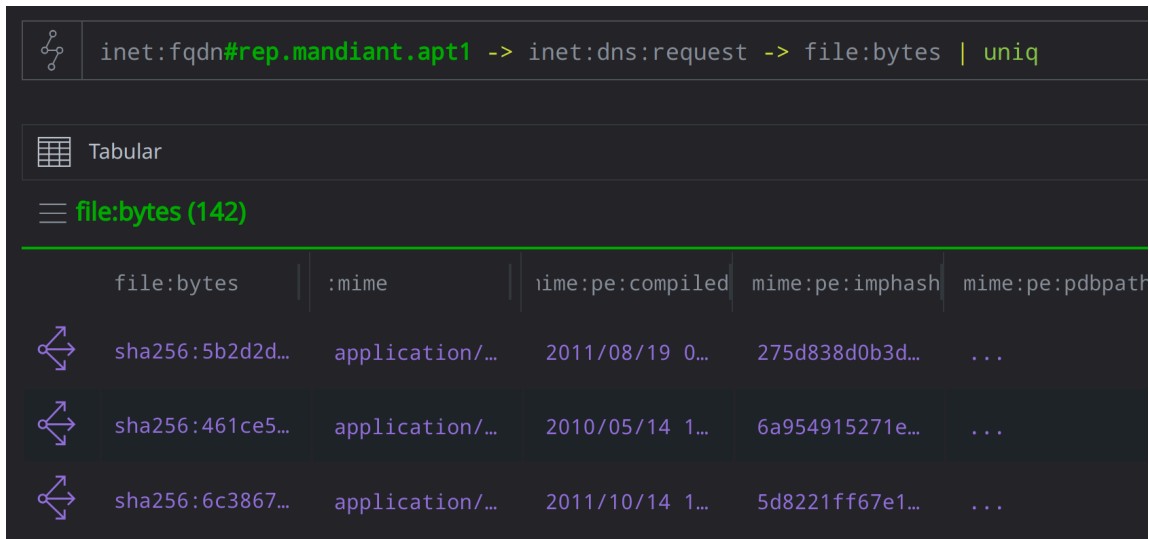
To answer our question, we're starting with **all** of the APT1 FQDNs - 2,073 of them.

When we **Explore**, we ask Synapse to take **all** 2,000+ nodes and show us **all** the things those nodes are connected to (over 54,000 nodes!) - even though we really only care about the DNS requests.

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

- When we use Storm to ask the **exact** question we want to answer, we get a response almost immediately!

There are **142** files that communicate with any of the 2,073 APT1 FQDNs:



```
inet:fqdn#rep.mandiant.ap1 -> inet:dns:request -> file:bytes | uniq
```

Tabular

file:bytes (142)

file:bytes	:mime	mime:pe:compiled	mime:pe:imphash	mime:pe:pdbpath
sha256:5b2d2d...	application/...	2011/08/19 0...	275d838d0b3d...	...
sha256:461ce5...	application/...	2010/05/14 1...	6a954915271e...	...
sha256:6c3867...	application/...	2011/10/14 1...	5d8221ff67e1...	...

When we ask Synapse to show us **precisely** the data we're looking for, Synapse is much more efficient - there is no need to select, display, and navigate through data we don't care about to (eventually) get our answer.