Hunting and Automation Using Open Source Tools

John Holowczak
Threat Researcher, Threat Analysis Unit
Carbon Black
@skipwich

Brian Baskin
Sr. Threat Researcher, Threat Analysis Unit
Carbon Black
@bbaskin

Carbon Black.

20 June 2019 - FIRST

CB TAU Threat Researchers

John Holowczak

- Development of TAU Research lab, Automation, Binary Analysis, Hunting
- Seasoned security operations analyst

Brian Baskin

- Research, Reverse Engineering, Automation, Hunting
- Formerly with US Defense Cyber Crime Center (DC3) focusing on military network intrusions

Carbon Black.

Agenda

Baselining: What and Why

Blueprinting your Organization

Automate the SOC

Threat Hunting

Baselining: What and Why

Carbon Black.

Know Your Environment

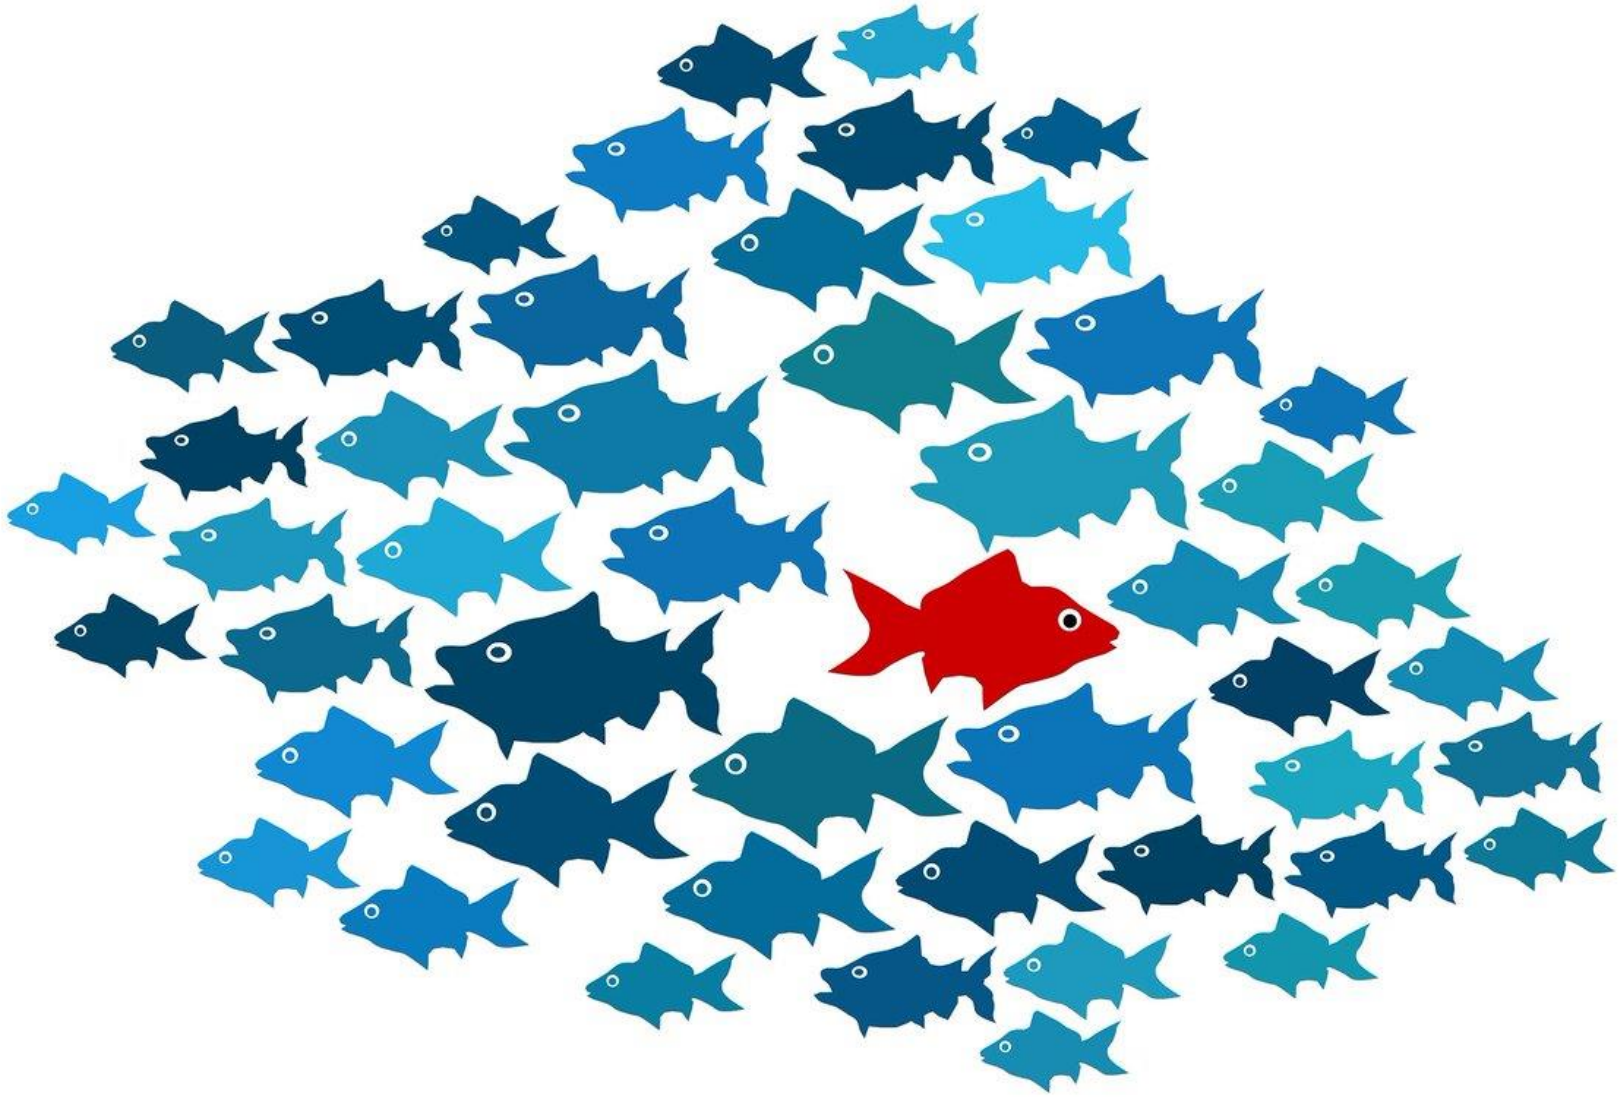

Carbon Black.

Know Your Environment

- With Baselining, turn over every stone; even normal behavior may be abnormal in reality
- When processing data, *classify* normal behavior and abnormal behavior
  - Certain behaviors can have multiple classifications
- Start your classification buckets large, add detail after each pass

Carbon Black.

Baselining vs Blueprinting Methods

Reactive

- Firehose of data
- Ingest all data into SIEM
- Tune False Positives Forever

Proactive

- Blueprint First
- Create rules for abnormal behavior
- Suffer less False Positives

Carbon Black.

Tools and Procedures

splunk>

sumologic

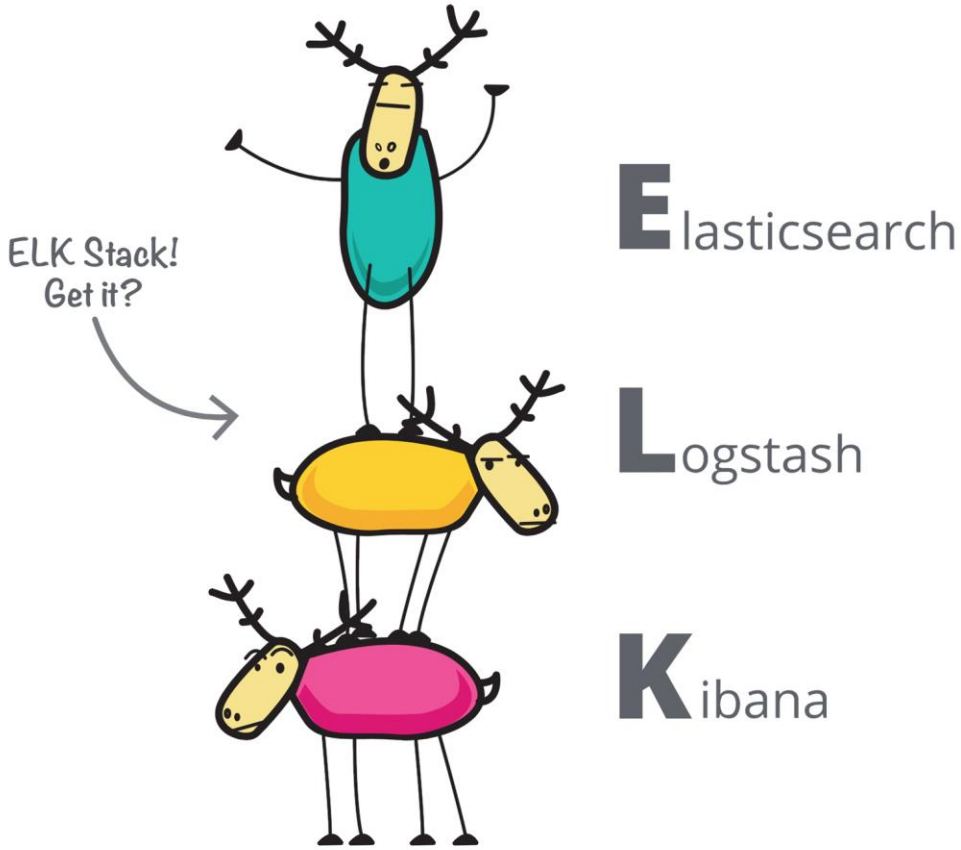

Carbon Black.

Enterprise Blueprinting

Carbon Black.

- Open source tool for querying endpoint metadata (at scale) like a database
- Utilizes SQL to expose data via a common interface
- Extensible in a number of languages
  - Add your own query-able data types

osquery

- Easy to get data from a number of endpoints at scale
- Quickly query data using a common language (SQL)
- Exhaustive list of metadata that is continually growing

- May be difficult to deploy across entire environment
  - Common orchestration tools can help with this (Ansible, Puppet, Chef)!

OSQuery Further Information

- List of schema available at https://osquery.io/schema/
- Some schema offer event information such as *process_file_events* which includes timestamps with when an event took place
  - Can only get this info if running OSquery in daemon mode, as it is an *evented table*
- Other file information schema:
  - Signature information
  - Startup items
  - Scheduled tasks

Carbon Black.

Low Prevalence Executables

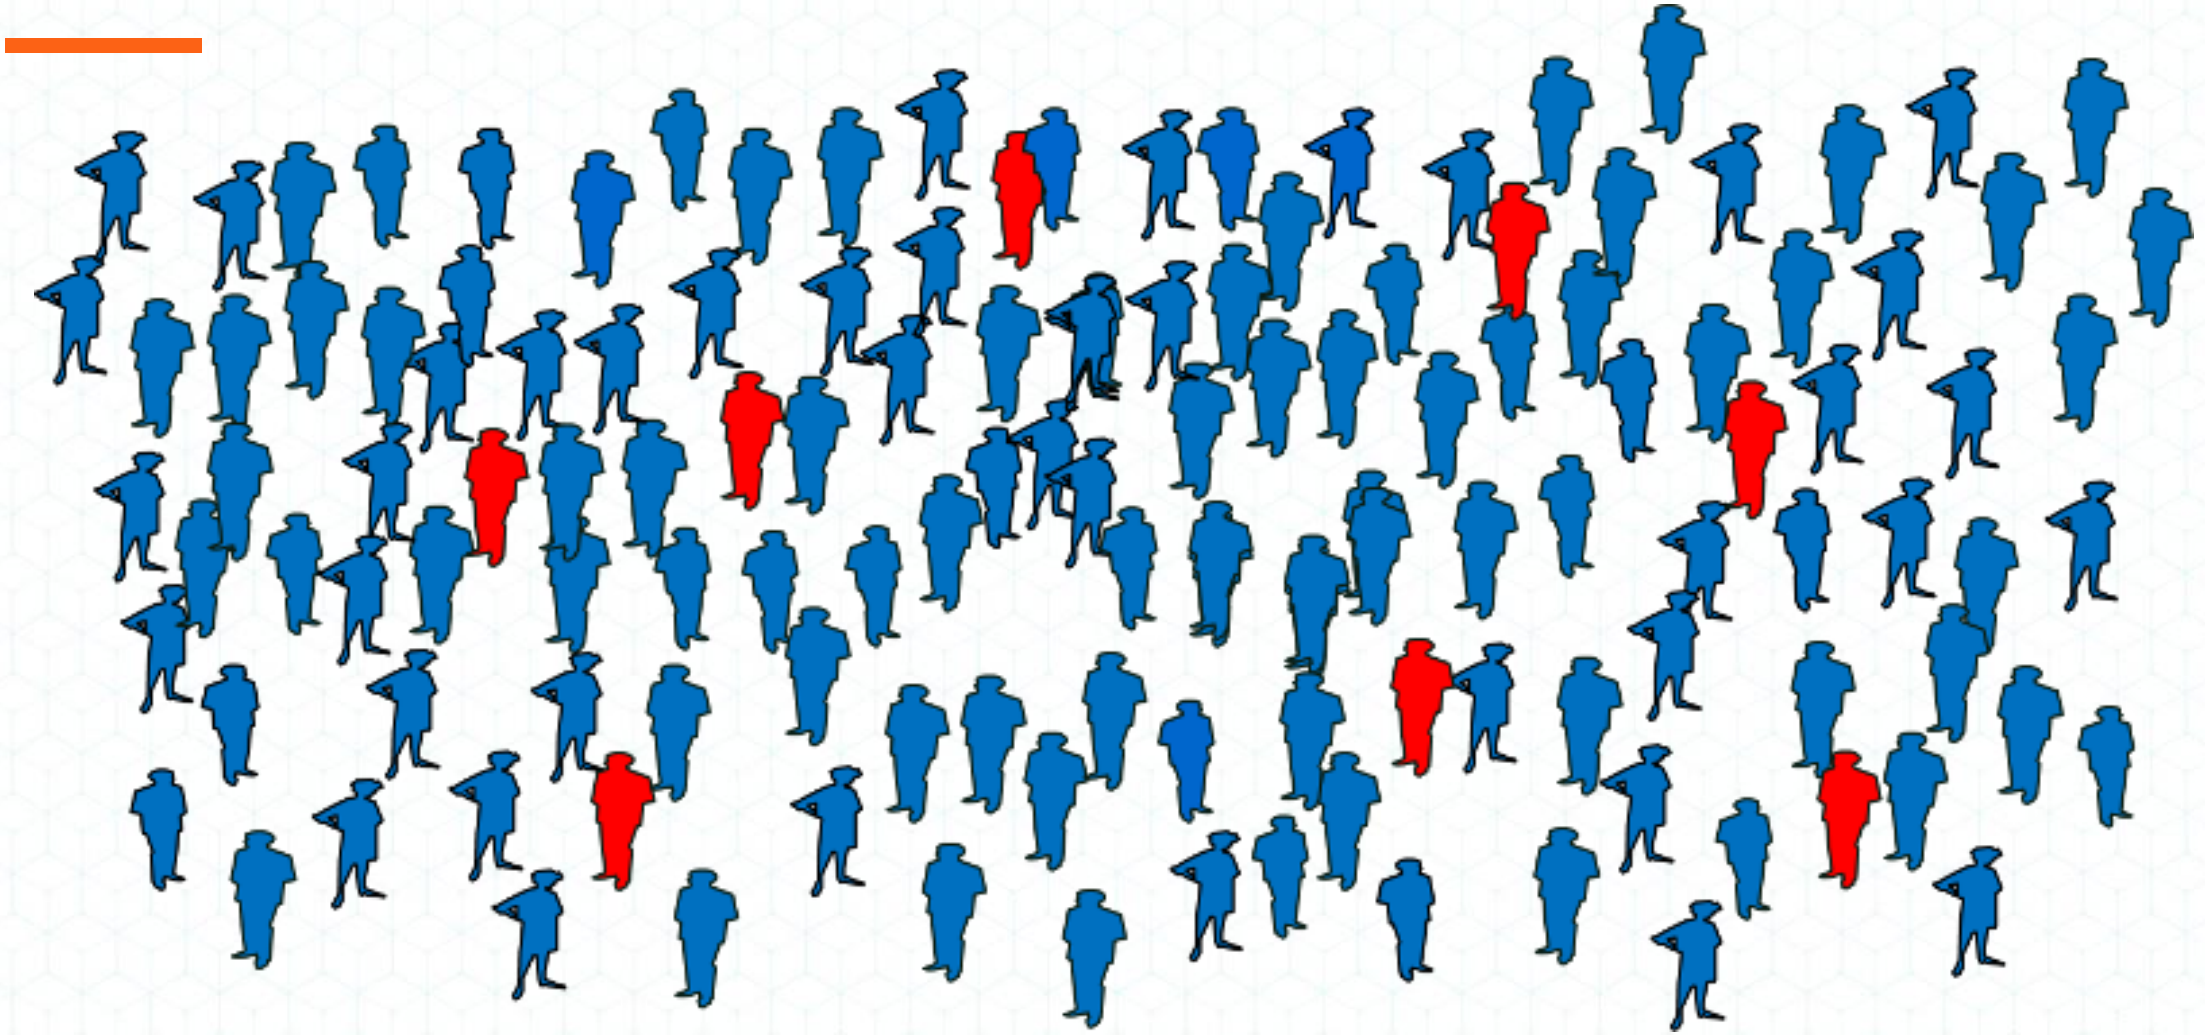

Carbon Black.

Low Prevalence Executables

- One-offs or rare applications
- Care less about the most common running programs
- Classify normal and abnormal for rarities to job functions


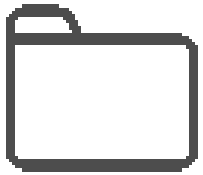


Carbon Black.

Leveraging OSQuery

```
osquery> SELECT name, pid, path, start_time FROM processes;
+----------------+-------+-----------------------------------------------------------------------------+------------+
| name           | pid   | path                                                                        | start_time |
+----------------+-------+-----------------------------------------------------------------------------+------------+
| systemd        | 1     | /usr/lib/systemd/systemd                                                    | 0          |
| rcu_sched      | 10    |                                                                             | 0          |
| migration/18   | 100   |                                                                             | 0          |
| udisksd        | 1002  | /usr/libexec/udisks2/udisksd                                                | 15         |
| ksoftirqd/18   | 101   |                                                                             | 0          |
| systemd-logind | 1024  | /usr/lib/systemd/systemd-logind                                             | 16         |
| gssproxy       | 1025  | /usr/sbin/gssproxy                                                          | 16         |
| irqbalance     | 1026  | /usr/sbin/irqbalance                                                        | 16         |
| smartd         | 1028  | /usr/sbin/smartd                                                            | 16         |
| kworker/18:0H  | 103   |                                                                             | 0          |
| lsmd           | 1032  | /usr/bin/lsmd                                                               | 16         |
| watchdog/19    | 104   |                                                                             | 0          |
| alsactl        | 1040  | /usr/sbin/alsactl                                                           | 16         |
| migration/19   | 105   |                                                                             | 0          |
| kworker/23:2   | 10555 |                                                                             | 17627181   |
| ksoftirqd/19   | 106   |                                                                             | 0          |
| mcelog         | 1068  | /usr/sbin/mcelog                                                            | 17         |
| kworker/19:0H  | 108   |                                                                             | 0          |
| kworker/16:2   | 10880 |                                                                             | 17627184   |
| watchdog/20    | 109   |                                                                             | 0          |
```


Carbon Black.

Networking Data

- SNMP (or equivalent) to pull data from your networking devices
- ARP Tables are a great start for network data collection
- Acquire IP and MAC addresses easily
- MAC Addresses are a great way to identify vendors on your network

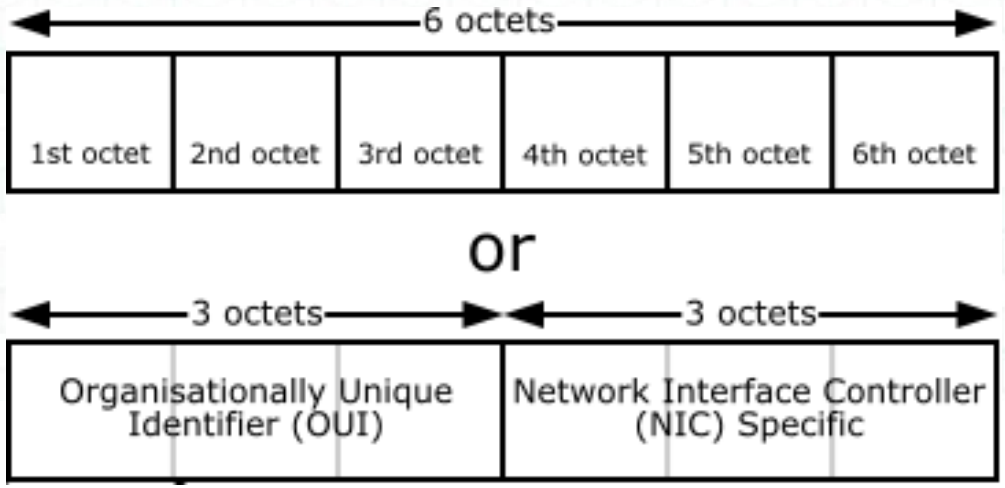

Common OUIs

https://www.wireshark.org/tools/oui-lookup.html

00:05:69 Vmware VMware, Inc.
00:0C:29 Vmware VMware, Inc.
00:1C:14 Vmware VMware, Inc.
00:50:56 Vmware VMware, Inc.

00:06:1B Notebook Notebook Development Lab. Lenovo Japan Ltd.
00:12:FE LenovoMo Lenovo Mobile Communication Technology Ltd.
00:59:07 Lenovoem LenovoEMC Products USA, LLC
0C:CB:85 Motorola Motorola Mobility LLC, a Lenovo Company
14:1A:A3 Motorola Motorola Mobility LLC, a Lenovo Company
14:30:C6 Motorola Motorola Mobility LLC, a Lenovo Company

00:00:97 DellEmc Dell EMC
00:01:44 DellEmc Dell EMC
00:06:5B Dell Dell Inc.
00:08:74 Dell Dell Inc.
00:0B:DB Dell Dell Inc.
00:0D:56 Dell Dell Inc.
00:0F:1F Dell Dell Inc.
00:11:43 Dell Dell Inc.
00:12:3F Dell Dell Inc.
00:12:48 DellEmc Dell EMC
00:13:72 Dell Dell Inc.
00:14:22 Dell Dell Inc.
00:15:30 DellEmc Dell EMC
00:15:C5 Dell Dell Inc.

00:60:B0 HP
08:00:09 HP
10:00:90 HP

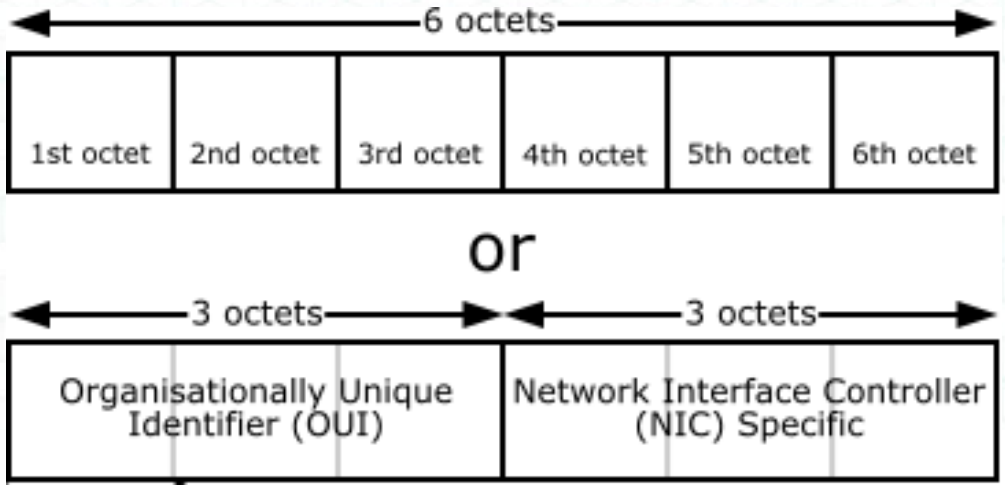

Using OSQuery to Enrich Networking Data

- OSQuery is a great tool to grab point-in-time endpoint data to supplement networking data
- Compare NICs and ARP tables on endpoint against Networking equipment ARP tables
- Great way to do full-coverage rogue detection

Carbon Black.

Getting ARP Data from OSQuery

- Using the *osqueryi* command locally we can test out our queries before running against whole environment

```
osquery> SELECT * from arp_cache;
+-------------+-------------------+-----------------+-----------+
| address     | mac               | interface       | permanent |
+-------------+-------------------+-----------------+-----------+
| 10.1        | 00:50:56          | em1             | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 10.1        | 00:0c:29          | em1             | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | br-423ba2418f06 | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 172.        | 02:42:ac          | docker0         | 0         |
| 10.1        | 00:00:00          | em1             | 0         |
| 10.1        | 00:50:56          | em1             | 0         |
```

```
osquery> SELECT a.address AS address, a.mac AS mac, a.interface AS interface, i.address AS interface_address
    ...> FROM arp_cache AS a
    ...> INNER JOIN interface_addresses AS i
    ...> ON a.interface = i.interface
    ...> WHERE a.interface = "em1";
+-------------+-------------------+-----------+-------------------------------+
| address     | mac               | interface | interface_address             |
+-------------+-------------------+-----------+-------------------------------+
| 10.         | 00:50:56:         | em1       | 10.                           |
| 10.         | 00:50:56:         | em1       | fe8                           |
| 10.         | 00:0c:29:         | em1       | 10.                           |
| 10.         | 00:0c:29:         | em1       | fe8                           |
| 10.         | 00:00:00:         | em1       | 10.                           |
| 10.         | 00:00:00:         | em1       | fe8                           |
| 10.         | 00:50:56:         | em1       | 10.                           |
| 10.         | 00:50:56:         | em1       | fe8                           |
| 10.         | 00:50:56:         | em1       | 10.                           |
| 10.         | 00:50:56:         | em1       | fe8                           |
| 10.         | 00:0c:29:         | em1       | 10.                           |
| 10.         | 00:0c:29:         | em1       | fe8                           |
| 10.         | 00:0c:29:         | em1       | 10.                           |
| 10.         | 00:0c:29:         | em1       | fe8                           |
| 10.         | 00:50:56:         | em1       | 10.                           |
```


Carbon Black.

SOC Automation

Easing the task of baselining

Carbon Black.

Automation Overview

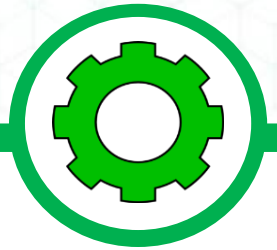

- Running minor, repetitive tasks
- Allows teams freedom to study data
  - Data Collection
  - Aggregation

- Scripts to find badness
  - This is what SIEM's do!
- Scripts making decisions for us
  - "This machine fell outside baseline, I'll automatically ban it"

Carbon Black.

Where do you put your data?

splunk>

sumologic

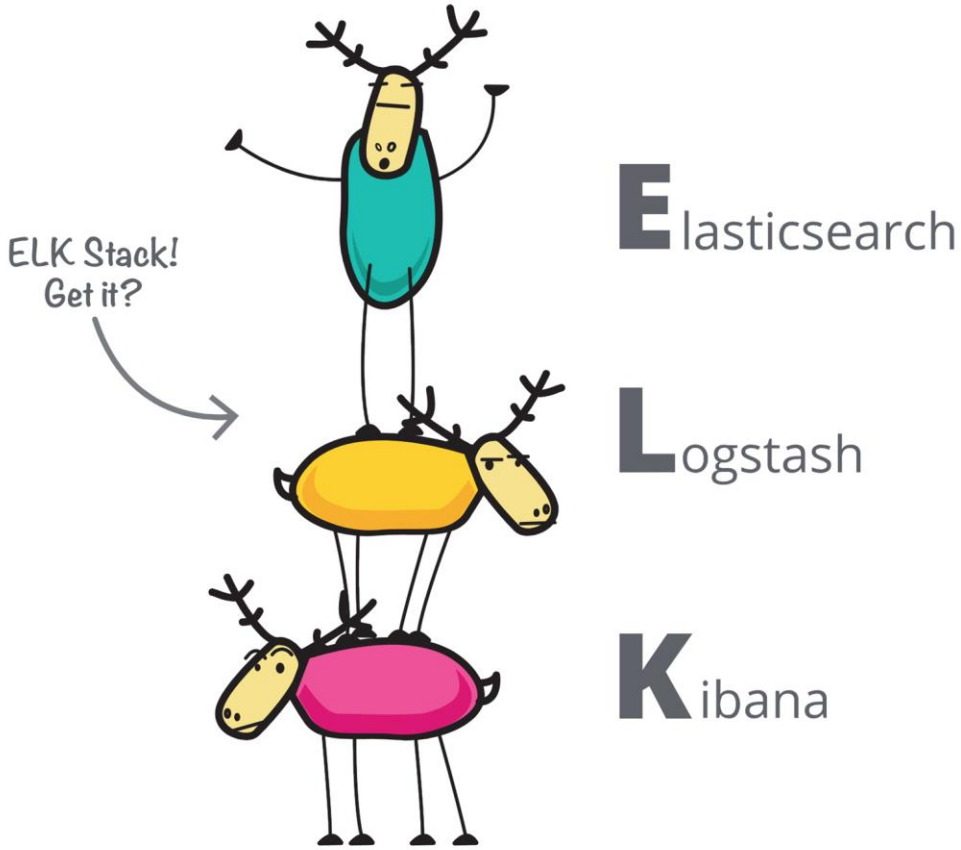

Carbon Black.

Data Collection

Querying Data

- Jupyter Notebook or R to automate studies
  - Programming heavy
  - Steeper learning curve
- ELK for querying
  - May take longer to set up
  - Large community support

Carbon Black.

Docker

- Public Docker for simple ELK stacks exist
- Require configuring data inputs (more involved)
- Seamlessly integrate OSQuery into Logstash using Filebeat (part of ELK)
- Easy-to-follow guide: https://elk-docker.readthedocs.io/

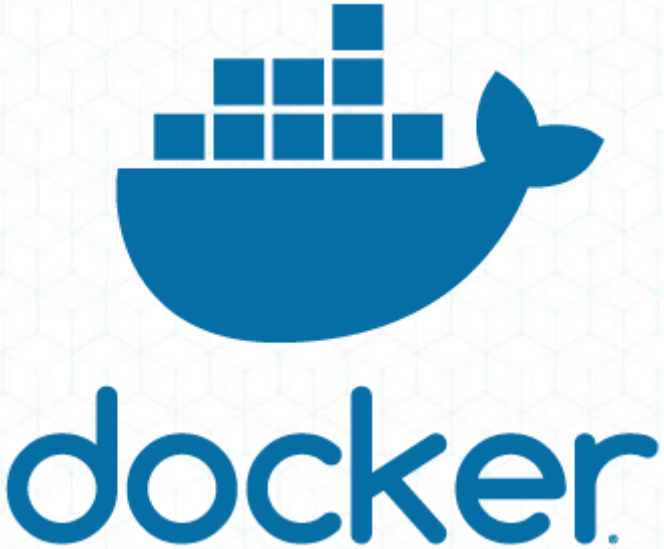

Filebeat

- Log forwarding service, part of ELK stack
- Has built-in templates for transforming OSQuery data into an easily-digestible format.
- OSQuery also has built in support for pushing to LogStash

Configuring OSQuery for Scheduled Queries

```
{
  "options": {
    "host_identifier": “hostname",
    "schedule_splay_percent": 10
  },
  "schedule": {
    "arp_cache": {
      "query": "SELECT * FROM arp_cache;",
      "interval": 10
    }
  }
}
```


Carbon Black.

Next Steps

- Forward OSQuery (using Filebeat) to Logstash, start hunting with Kibana

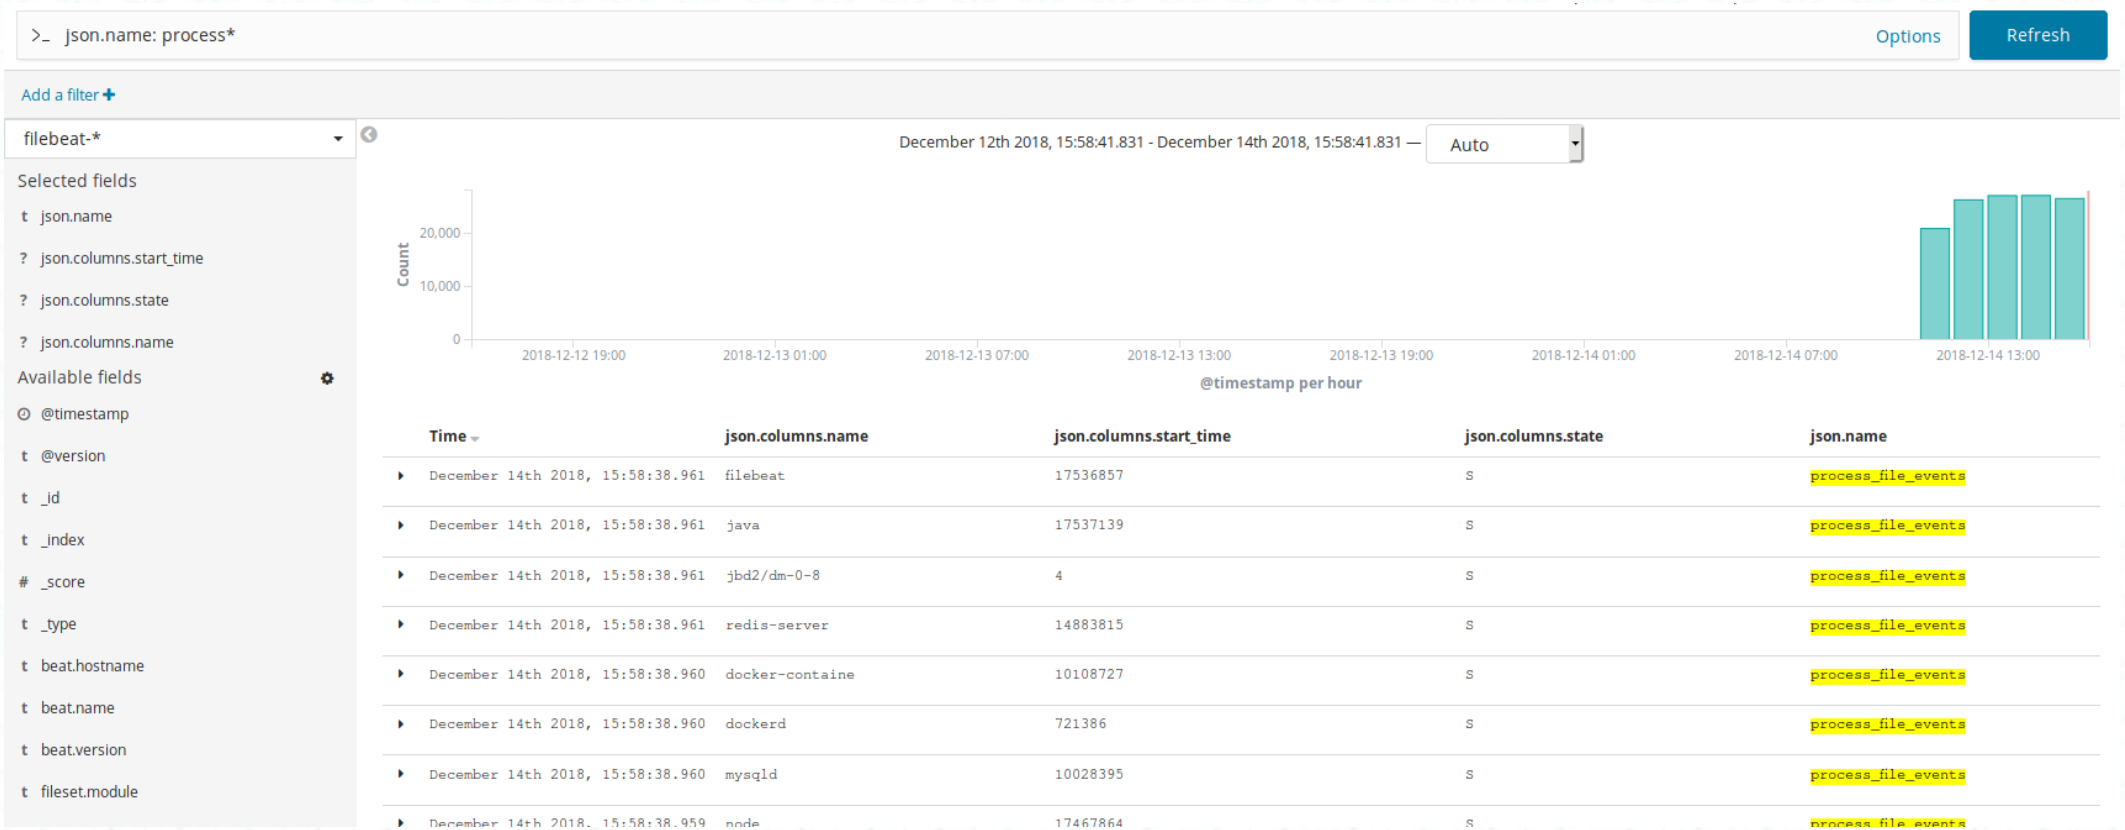

Carbon Black.

Threat Hunting

Tying it all together

Carbon Black.

Using Statistical Analysis for Threat Hunting

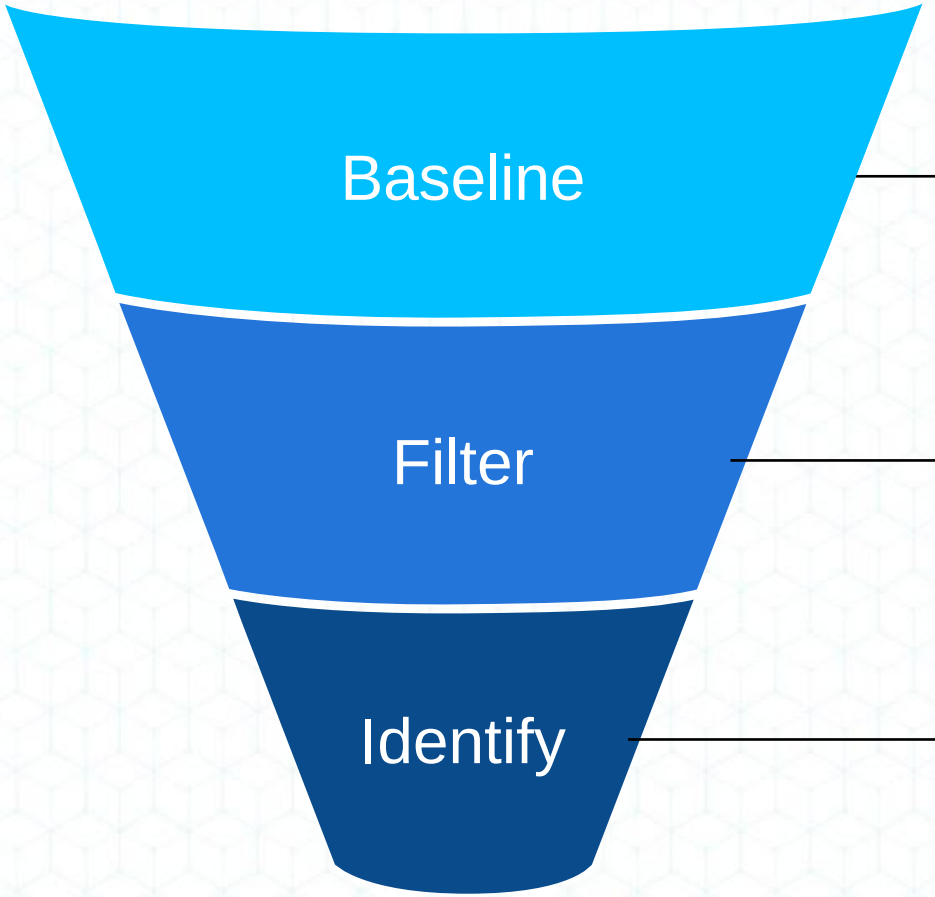

- Determine the “normal” dataset
- Eliminate baseline as noise
- Investigate anomalies & statistical deviations

Carbon Black.

Analyzing Data

In [6]:

```
sns.set(style="white", palette="muted", color_codes=True)
f, axes = plt.subplots(1, 1, figsize=(7, 7), sharex=True)
sns.despine(left=True)
sns.distplot(cs, hist=False, color="g", kde_kws={"shade": True})
plt.setp(axes, yticks=[])
plt.tight_layout()
```

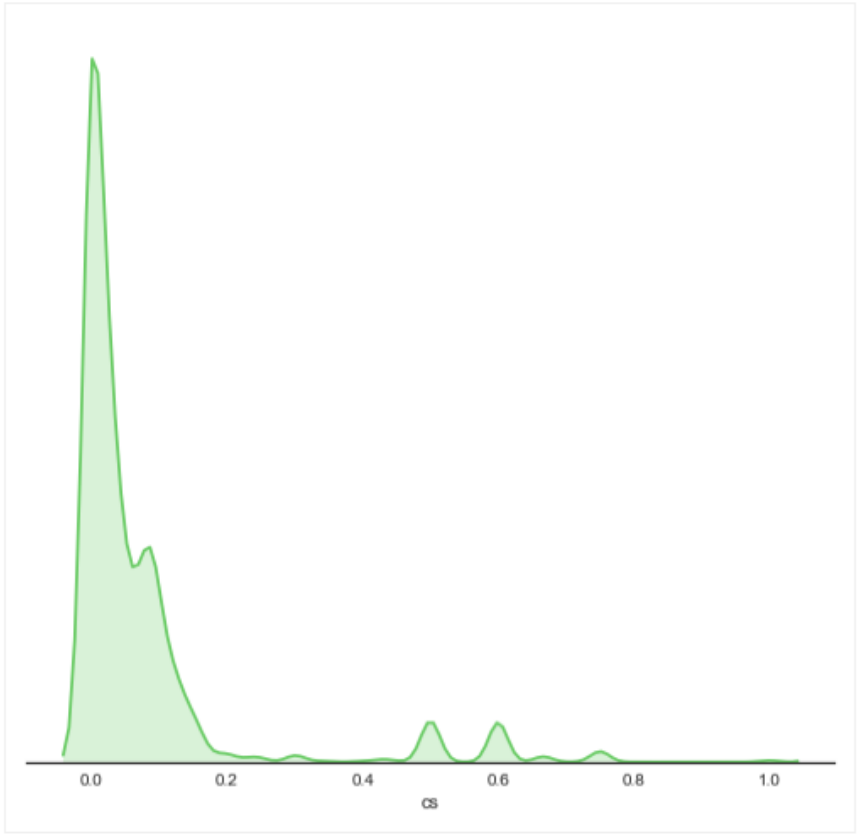

Carbon Black.

Hunting Methodologies

- Back to the basics: Now time to look for the abnormal
- Search across environments for behavior and static IOC's
- Least prevalent occurrences tend to be most abnormal

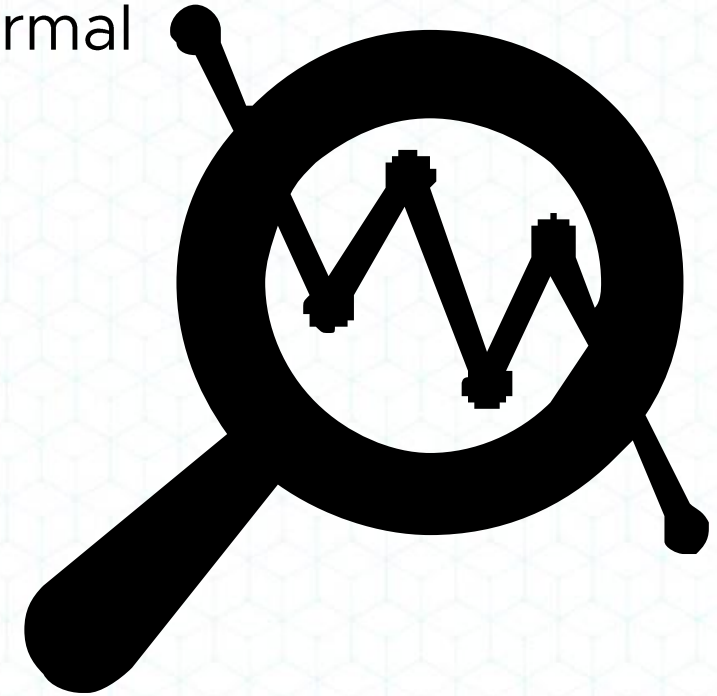

Carbon Black.

MAC Addresses – Uncommon Environment OUIs

10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	52:54:00	RealtekU	Realtek (UpTech? also reported)
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.
10	00:50:56	Vmware	VMware Inc.

Carbon Black.

OSQuery Hunting

```
osquery> SELECT name, pid, path, start_time FROM processes;
+----------------+-------+----------------------------------------------------------------------------+------------+
| name           | pid   | path                                                                       | start_time |
+----------------+-------+----------------------------------------------------------------------------+------------+
| systemd        | 1     | /usr/lib/systemd/systemd                                                   | 0          |
| rcu_sched      | 10    |                                                                            | 0          |
| migration/18   | 100   |                                                                            | 0          |
| udisksd        | 1002  | /usr/libexec/udisks2/udisksd                                               | 15         |
| ksoftirqd/18   | 101   |                                                                            | 0          |
| systemd-logind | 1024  | /usr/lib/systemd/systemd-logind                                            | 16         |
| gssproxy       | 1025  | /usr/sbin/gssproxy                                                         | 16         |
| irqbalance     | 1026  | /usr/sbin/irqbalance                                                       | 16         |
| smartd         | 1028  | /usr/sbin/smartd                                                           | 16         |
| kworker/18:0H  | 103   |                                                                            | 0          |
| lsmd           | 1032  | /usr/bin/lsmd                                                              | 16         |
| watchdog/19    | 104   |                                                                            | 0          |
| alsactl        | 1040  | /usr/sbin/alsactl                                                          | 16         |
| migration/19   | 105   |                                                                            | 0          |
| kworker/23:2   | 10555 |                                                                            | 17627181   |
| ksoftirqd/19   | 106   |                                                                            | 0          |
| mcelog         | 1068  | /usr/sbin/mcelog                                                           | 17         |
| kworker/19:0H  | 108   |                                                                            | 0          |
| kworker/16:2   | 10880 |                                                                            | 17627184   |
| watchdog/20    | 109   |                                                                            | 0          |
```

OSQuery Hunting

```
SELECT name FROM processes WHERE start_time < 100;
```

```
   7189 csrss.exe
   7189 lsass.exe
   7189 conhost.exe
    ...
   2348 dllhost.exe
    103 firstboot.cmd
     45 FlashPlayerUpdateService.exe
      8 psexesvc.exe
      3 rund11.exe
      1 conh0st.exe
```

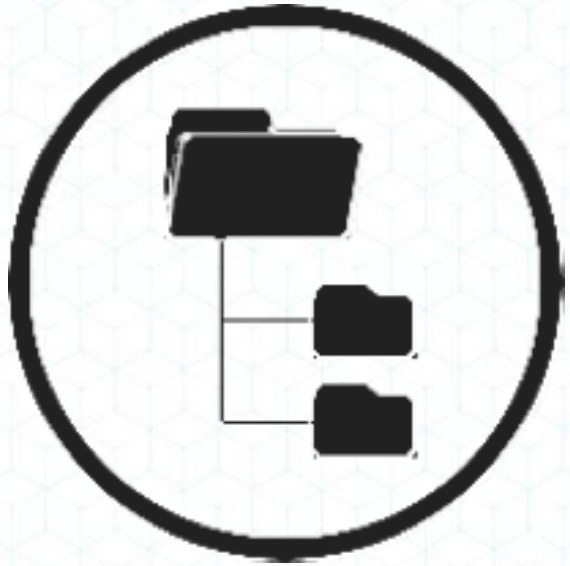

Prevalence of Executables

- Can you:
  - Identify abnormal software running on fewest endpoints?
  - Identify executables that are widespread but in unusual places?

- Yes!
  - Extract data on binaries from osquery
  - Combine into CSVs and perform text magic

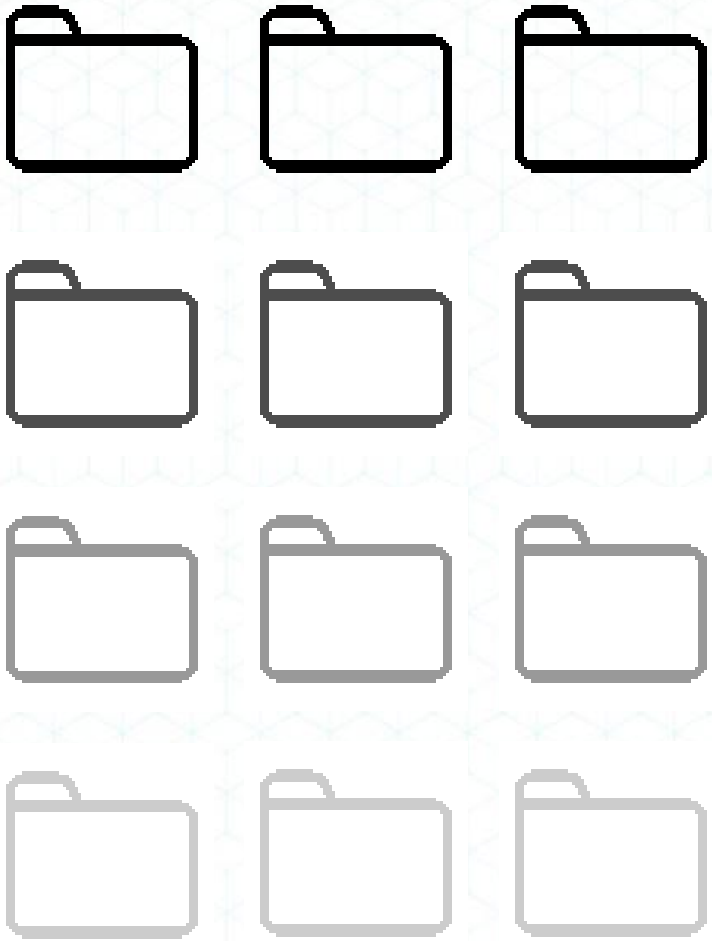

Filtering Data

- Expressions to hunt for unusual indicators
  - Files that have a single character filename:
    \\.\....,
  - Files running one-folder deep from volume root:
    (:\\[a-zA-Z0-9]{1,12}\\[a-zA-Z0-9]*\....,)
  - Files run directly from Windows folder:
    (:\\windows\\.{1,15},)
  - Files with unusual extensions:
    (\.bin,|\.dat,|\.log,|\.gif,|\.txt,|\.jpg,|\.rar,|\.sql,)

Mass Searching

- One-character file names:

```
   6 c:\tdm-gcc-64_4.9.2\work\a.exe
   1 c:\accbk\agusta\y.bat
   1 c:\users\jsmith\appdata\local\microsoft\windows\temporary internet
files\content.ie5\4unu162n\..exe
   1 sysvol\users\z9service\downloads\q.exe
   1 sysvol\program files (x86)\k2 for sharepoint 2013\z.bat
```

- Low prevalence in Windows Folder

```
  22 c:\windows\psexesvc.exe
   1 c:\windows\system32\oem\firstboot.cmd
   1 sysvol\windows\system32\dsget.exe
   1 c:\windows\system32\hpbpro.exe
   1 c:\windows\system32\scardsvr.exe
```

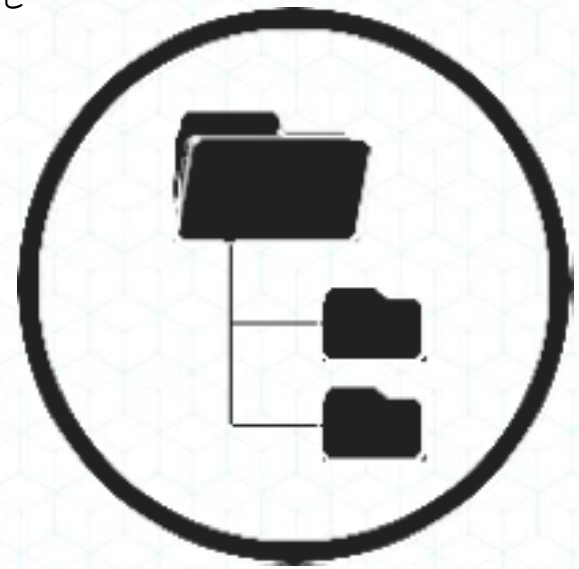

Carbon Black.

A Story of Two Executables (PLink)

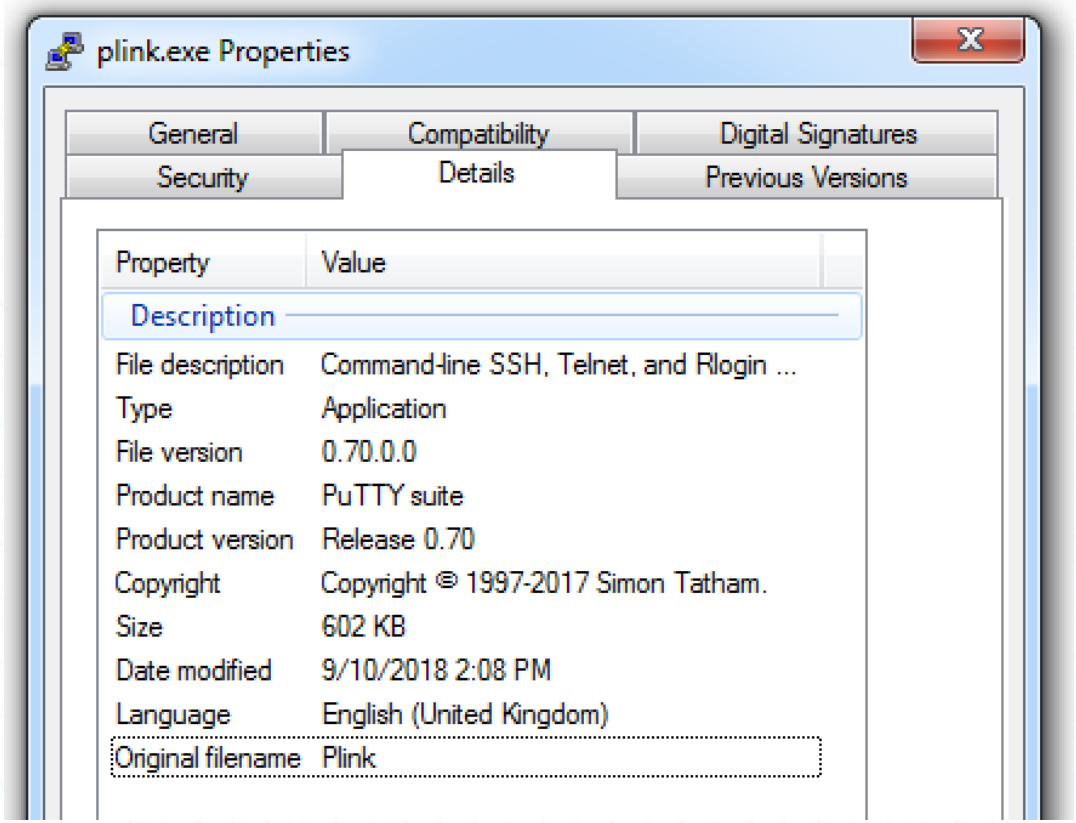

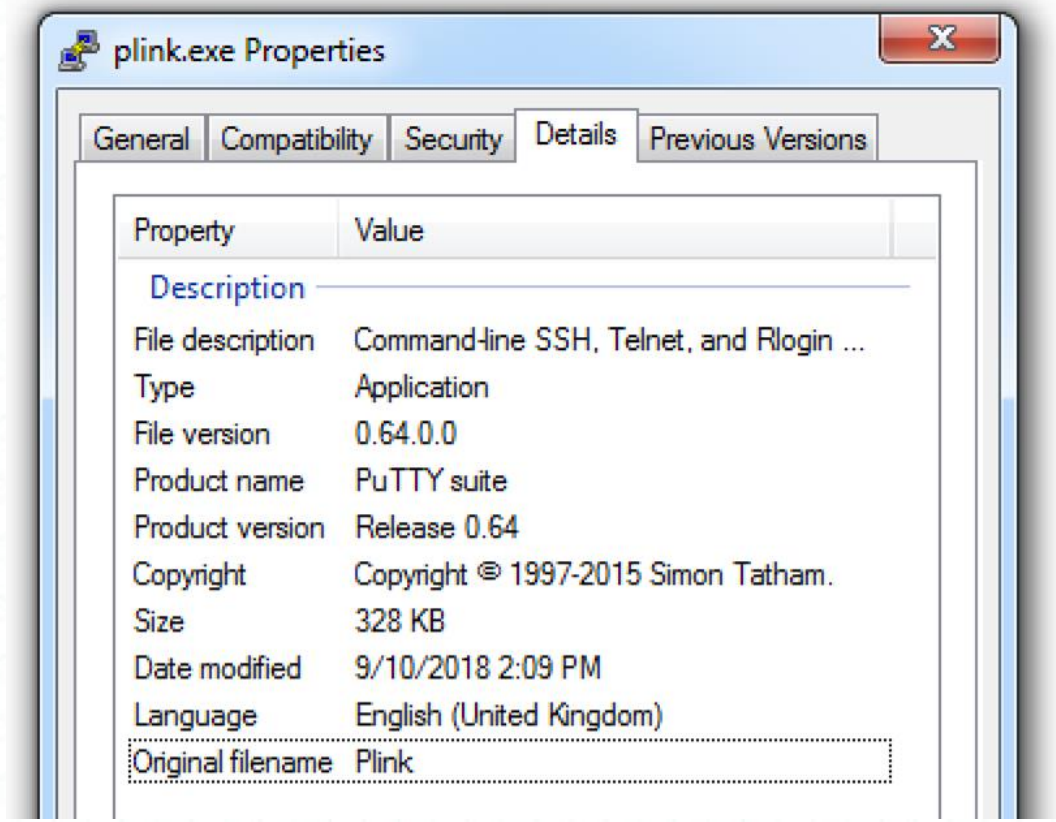

Carbon Black.


Happy Hunting!
Carbon Black.

John Holowczak
@skipwich
jholowczak@carbonblack.com

Brian Baskin
@bbaskin
bbaskin@carbonblack.com

Carbon Black.

20 June 2019 - FIRST