gravwell

CASE STUDY

GRAVWELL IN THE SCINET NOC

An action report and case study in threat hunting

SciNet SC18 Dallas, TX | hpc inspires.

Table of content

EXECUTIVE SUMMARY

For the 2018 SC Conference (SC18, held in Dallas, TX), Gravwell provided our analytics platform to the Network Security team.

These brave souls were responsible for cyber security on a network consisting of $52 million in contributed hardware, software, and services plus 4.02 Terabits per second of external capacity. This means that not only does the SCinet Network Security team need to protect SCinet from the world, it needs to protect the world from SCinet.

This is a challenging task but we were excited to give it a go and I think the results were spectacular. Jason Zurawski, SCinet chair for the conference, observed "The SCinet is purposely designed to facilitate experimentation for new hardware, software, and services. We are pleased to support emerging companies, such as Gravwell, as they pioneer new products and learn from performance of our network and the experience of our volunteers."

And learn we did! We learned that Gravwell is not only up to the task of handling these kinds of analytics, but we also did it on significantly less hardware than previous years. During the event, Gravwell ingested over 4.6 billion entries comprising over 1TB of data from a variety of sources. Analysts ran 4281 manual searches, 17325 automated searches, and viewed dashboards 1159 times during the two weeks in the Network Operations Center (NOC).

All those numbers seem great but what was the actual impact for the team? The SCinet Network Security team benefited in two major ways. First, a good chunk of tedious analysis and investigation was automated with Gravwell which freed up analysts to focus on threats that mattered. Secondly, investigations were expedited using Gravwell pre-built investigation dashboards and since insights are built off of actual data, not metadata translations, root-cause analysis is always possible.

At the event, the SCinet Network Security team used Gravwell to stop continuous internet attacks automatically. With a good chunk of busy work removed, the team was freed up to better to identify, hunt, and respond to an actual attack that sought to bring the entire force of 4.02 Tb/s against an unsuspecting SaaS company. Thanks to a crack team and the power of Gravwell, the day was saved.

Book some time with the Gravwell team to implement this level of defense in your organization by emailing sales@gravwell.io or visiting https://www.gravwell.io/schedule-a-demo.

Keep reading for detailed information about the event, the Network Security Team, and to follow along with the threat hunt.

BACKGROUND

What is SkiNet?

SCinet is the SC Conference's dedicated high-capacity network infrastructure, designed and built by volunteer experts from industry, academia, and government.

Planning begins more than a year in advance of each SC Conference and culminates in a high-intensity installation that, for the duration of the conference, is the fastest and most powerful network in the world.

SCinet gave attendees the chance to experience t he world's fastest temporary network, delivering 4.02 terabits per second of wide area capacity to the Kay Bailey Hutchison Convention Center Dallas.

In preparation volunteers installed more than 67 miles of fiber optic cable, including two miles of new underground fiber that now connects the convention center to a downtown Dallas data center. After the conclusion of this year's conference, that underground fiber remains in place for the benefit of the city of Dallas.

To deliver WiFi for all attendees across one million square feet of exhibit space, volunteers also installed 300 wireless access points in just one week.

Here's the SCinet overview video put out by the team:

www .youtube .com/ watch?v=B26DCSCl-7Q

gravwell

CASE STUDY

ARCHITECTURE

SCinet is made possible by the contributions of 40 industry-leading organizations

Who in total donated $52 million in hardware, software, and services

As you can imagine, handling security on this type of network has many challenges.

gravwell

CASE STUDY

SCinet Network Security

THE JOB

The job of the SCinet Network Security team is to minimize malicious activity on the SCinet network

And to provide as safe a haven as possible for the SC attendees, exhibitors, researchers, and organizers. As such, they need to protect the SCinet infrastructure from the internet, but also to protect the internet from SCinet.

Primary focus areas are to conduct vulnerability assessments of SCinet infrastructure, incorporate threat intel, monitor for threats and alerts, and mitigate where needed. We'll be focusing on the monitoring aspects because that's where our case study takes place.

THREATS AND ALERTS

The team focuses on critical infrastructure that may impact services and those compromises that can affect the experience. Some threats (e.g. DMCA complaints, incidents reported from booths and other teams) come to our attention from outside SCinet Network Security, but most threats are identified by Network Security vendors, tools, and team members' analyses. The principle of "First do no harm" aka "Don't be an agent of DOS" is used to moderate security response to potential threats.

Nuisance behavior is not in and of itself sufficient reason to disable access for SCinet attendees; packets happen. For example, vulnerable exhibitor or attendee hosts on the SCinet network do not generally present a threat to SCinet, though they can subsequently become compromised and engage in clear malicious activity. Vulnerabilities and suspected nuisance activity are worthy of contacting the user and offering assistance, though this is only usually feasible for eduroam and booth services, and this is a lower priority than mitigation of bona fide malicious activity.

DATA SOURCES

Gravwell ingested data from many sources over the course of the conference. In addition to a ton of host-based syslog, we also collected logs from network security appliances.

The Reservoir Labs R-Scope products provided a huge volume of Bro-formatted logs down to the level of individual connections across the network. Attivo's BotSink product stood up decoy virtual machines and sent in logs about attempted attacks.

DATA ENRICHMENT

For the event, the SCinet Network Security team made use of some open source data enrichment and threat feed capabilities. For threat detection we were using malware domains dns blacklist and virustotal. Integrating the threat feed allowed us to monitor for known threat actor activity in the network throughout the event with automated DNS auditing.

We also utilized the Maxmind IP geolocation database for layer3 and MAC->manufacturer resolution for layer2 traffic analysis.

In addition to generic sources we were enriching via hostname lookup, VLAN naming, infrastructure details, and other organizationally specific information.

SOUND INTERESTING?

Check out https://www.gravwell.io/blog/auditing-dns-with-coredns-and-gravwell.

ANALYSIS TECHNIQUES

A good portion of the analysis being conducted in the SCinet NOC was done autonomously;

We utilized Gravwell scheduled searches to create an autonomous SOC/NOC that conducted basic threat hunting and tip confirmation.

EXAMPLE AUTONOMOUS SOC/ NOC ACTIVITY

For the conference, we implemented a number of autonomous operations in order to free up resources for active hunting and provide automatic threat blocking where confidence levels were high enough.

The SCinet Network Security team incorporated Attivo "network based threat deception" decoy systems into the infrastructure to provide detection and threat intelligence on any attacker activity against those systems (https://attivonetworks.com/product/attivo-botsink/). These devices fed logs into Gravwell.

One of the autonomous activities we created was to monitor the Attivo logs for brute force SSH activity.

The Attivo decoys were configured to allow an attacker entry after a dynamic number of failed login attempts. The results of the search would show any IP address attempting to gain unauthorized access to the systems. Gravwell can go beyond just detection and reporting of this type of activity.

Thus, utilizing a combination of technologies, we could customize automation of tedious SCinet Network Security work to our needs and free up valuable analyst time to work on more dynamic and challenging problems. This one example of automating Gravwell + Attivo resulting in blocking hundreds of IPs and saved our analysts valuable time. Instead of chasing down "script kiddie" activities like bots and brute forcing, SCinet Network Security team members could focus on the threats that actually mattered.

FINDINGS

The team created a variety of dashboards to monitor activity as the event progressed . As is tradition for security teams at computing conferences, there's a "wall of sheep" dashboard that covers the low hanging fruit for attackers . This would be things like passwords submitted over HTTP instead of HTTPS, telnet activity, etc . We also included results from the Attivo decoys .

One of the student volunteers was quite interested in some other more social analytics such as which dating app was most popular, which we analyzed using mostly DNS traffic .

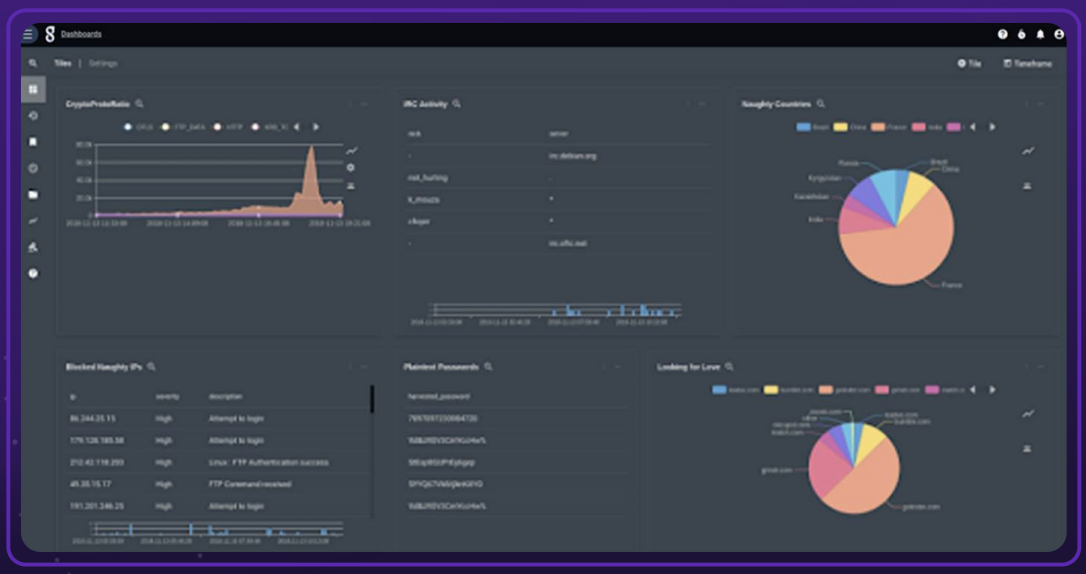

We also created some overview dashboards to monitor general network and infrastructure activity:

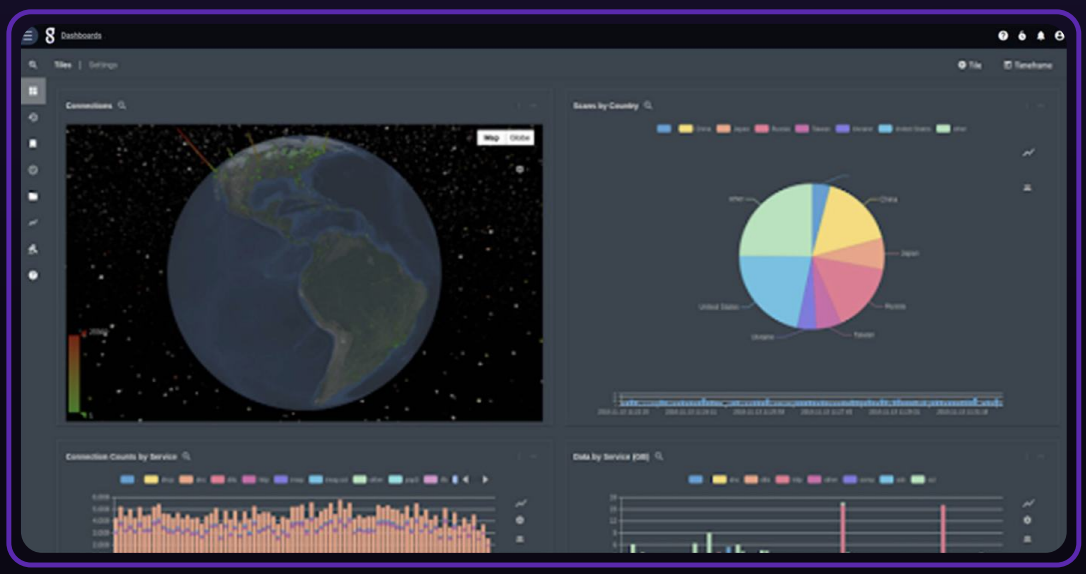

There were a number of investigations conducted and one of them stood out as a textbook case study for hunting activity for a few reasons.

The remainder of this section covers that example.

CASE STUDY

Hunting adversaries with Gravwell

This is a redacted write-up of the hunt we did on a successful attack that occurred on Nov 15th. In summary, an attacker gained a foothold during initial setup of a vendor system that was part of the SCinet infrastructure due to an easily brute forceable password. As part of the setup process, the vendor properly changed the default password which resulted in no vulnerabilities being found during weak password assessment by the Network Security team.

Thus, the SCinet Network Security team was not aware of a potential issue until the dormant malware came alive nearly two weeks after initial compromise.

The bulk of the investigation was conducted utilizing Bro logs generated by Reservoir Labs' R-Scope. We also included data sources such as IPFIX, system logs, and data enrichment like a DNS blacklist from http://www.malwaredomains.com/.

The conference requested the redaction of any and all IP addresses for this report .

To provide contextual clarification, SCinet IP addresses have been colored green .

THE TIP

InMon (https://inmon.com/) was observing the switches and providing an aggregate bandwidth dashboard. An operator noticed an uptick in traffic on port 22466 on the morning of 11/15/2018. As is often the case, attackers try to mask themselves in the noise of daily operations. This happened to occur on the day that the SCInet bandwidth test is conducted –when massive amounts of network traffic are sent on purpose in order to test the throughput.

However, this anomaly started prior to the designated start time of that test and so the operator reported the tip. We had set up IPFIX ingestion directly from networking equipment earlier in the week so we used that data feed to confirm the tip. We could have used the “conn-long” Bro logs generated by Reservoir Labs’ R-Scope but the binary nature of IPFIX makes it faster for this search. We confirmed the InMon tip data with maxmind enrichment to note suspect behavior:

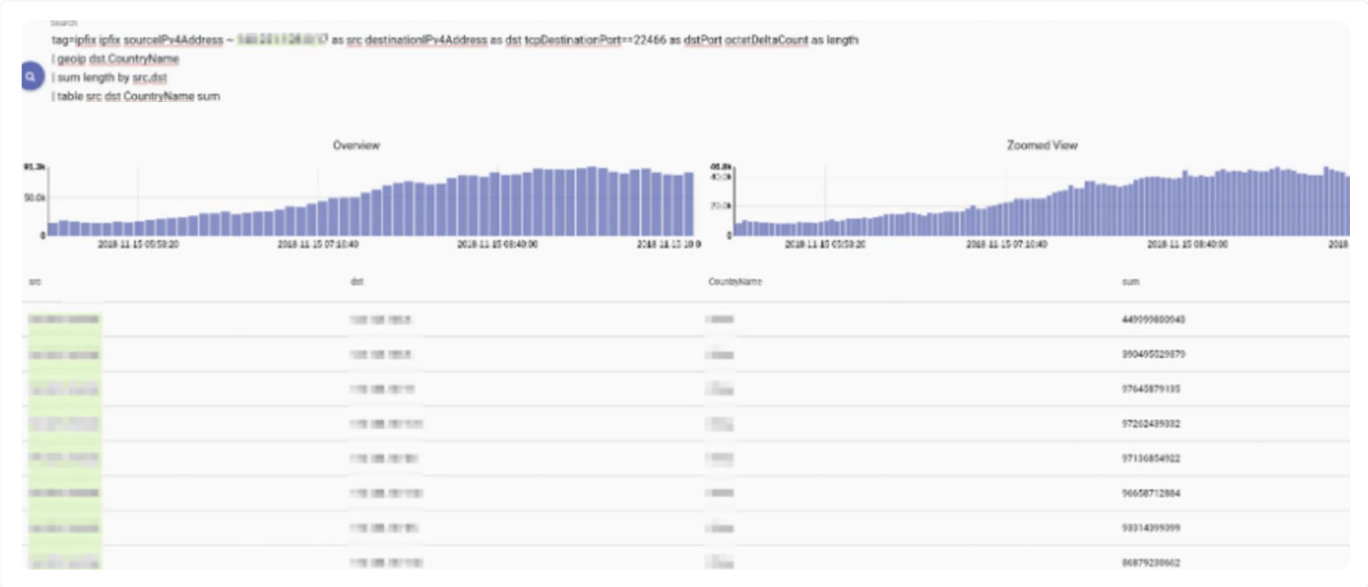


```
tag=ipfix ipfix sourceIPv4Address ~ xxx.xxx.xxx.0/17 as src
destinationIPv4Address as dst tcpDestinationPort==22466 as dstPort
octetDeltaCount as length

| geoip dst.CountryName
| sum length by src,dst
| table src dst CountryName sum
```

THE HUNT

An investigation was started into the offending IP addresses of xxx.xxx.xxx.xx1
and xxx.xxx.xxx.xx2 which were transmitting large amounts of data over port 22466 to an overseas IP address.

As part of the activity for SCinet, we used our “IP Investigation Dashboard” which contains a bunch of pre-built searches to do hostname lookups, DHCP enrichment, show DNS activity, HTTP requests, geolocation maps, and much more. Basically the first steps for investigating a suspicious IP.

There are some interesting domains and URLs here:

```
ppp.gggatat456.com
www1.gggatat456.com
ppp.gggatat456.com
navicatadvvr.com
wowapplecar.com
navicatadvvr.com
ppp.xxxatat456.com
wowapplecar.com
ppp.xxxatat456.com
navicatadvvr.com
topbannersun.com
navicatadvvr.com
ppp.xxxatat456.com
xxx.xxx.xxx.xxx/c.txt
```

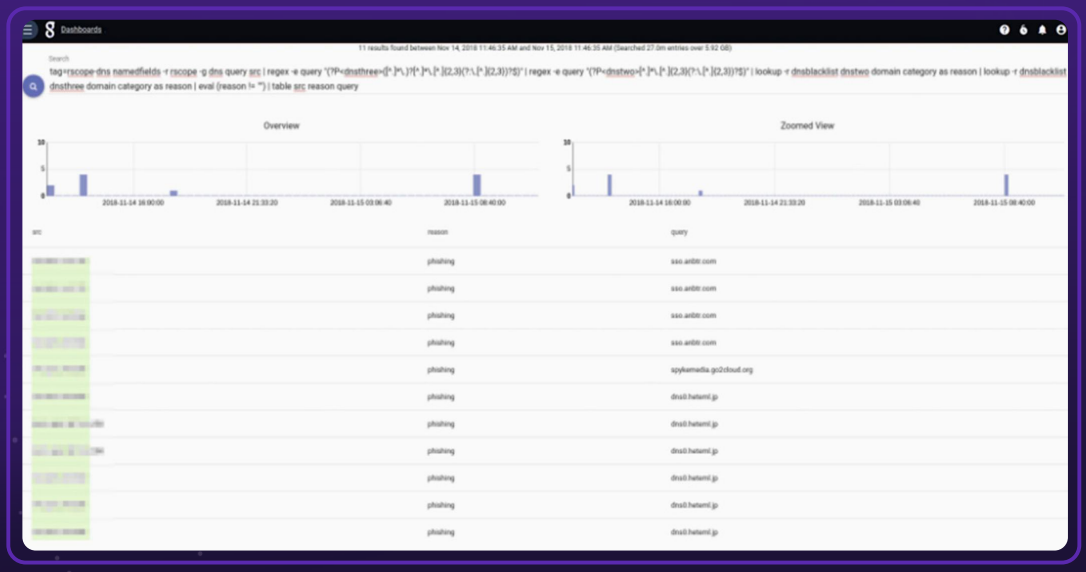

We used the malwaredomains blacklist to act as a DNS tip from the DNS logs extracted from the R-Scope systems

but in this instance we had no matches. So this activity isn't yet reported or well known, or is otherwise custom for this threat actor.

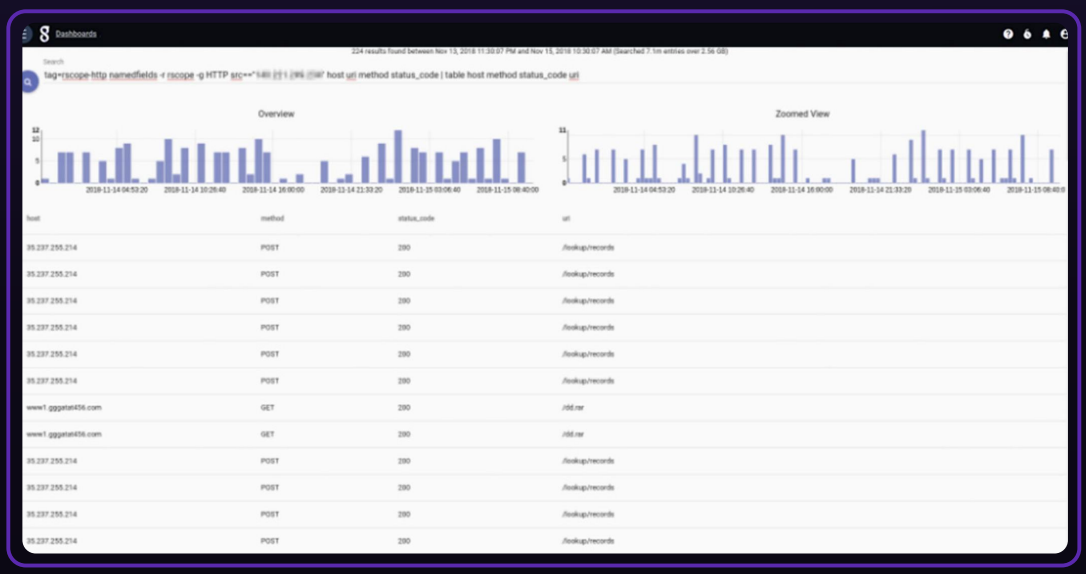

One of the things that we noticed were some suspicious http activity originating from the suspect hosts to a few domains and URLs.

We needed to further investigate this search so we expanded the http tile for more detailed information.

I reached out to grab the payload from the server using wget:

Looks like an encrypted blob (entropy is 5.975220 bits per byte) which is not at all unexpected. It likely contains instructions for the bot to execute. Attempting to resolve many of the domains was proving fruitless but a couple of requests were made to a direct IP

```
remasis@michelangelo:~/tmp/scinet/malware$ md5sum dd
5fc38cdccd16710969b09582c9af34dc  dd.rar
remasis@michelangelo:~/tmp/scinet/malware$ file dd.ra
dd.rar: data
remasis@michelangelo:~/tmp/scinet/malware$ xxd dd.rar
00000000: 2f21 077b 4c3e 5224 2c38 5045 0904 6803  /!.{L>R$,8PE..h
00000010: 7776 1c73 711a 0470 756d 0b04 061f 7700  wv.sq..pum....w.
00000020: 7b6b 0071 6f00 0f6d 7070 0b1b 0503 7f1c  {k.qo..mpp....
00000030: 7671 1c77 7403 3b4b 2428 5550 5a50 2b57  vq.wt.;K$(UPZP+W
00000040: 7f6a 5732 221b 4e39 3a39 434d 181e 3e4a  .jW2".N9:9CM..>J
00000050: 3a48 3834 2c52 5f2d 277c 1641 5941 691c  :H84,R_-'|.AYAi.
00000060: 3136 5a77 7318 1935 2f31 161b 4742 2e00  16Zws../1..GB..
00000070: 7145                                     qE
remasis@michelangelo:~/tmp/scinet/malware$
```

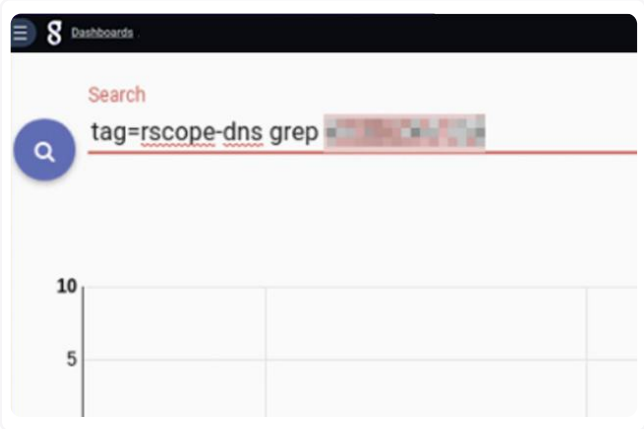

This IP address was almost certainly given to the compromised host via the C&C blob but just to verify that hypothesis, let's run a basic search to see if anyone requested a domain name that resulted in that address:

```
tag=rscope-dns grep xxx.xxx.xxx.xxx.
```

Hypothesis confirmed as we see a return of 0 DNS answers with that IP address.

We grabbed the payloads for cursory analysis:

In my not-so-novice opinion, a file called 'c.txt' that's actually an ELF binary is bad, mmmmkay. Quick submission to virustotal and we've got easy confirmation:

```
remasis@michelangelo:~/tmp/scinet/malware$ wget [redacted]/c.txt
--2018-11-15 10:49:37--  http://[redacted]/c.txt
Connecting to [redacted]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 625611 (611K) [text/plain]
Saving to: 'c.txt'

c.txt                                              100%[==========>]
==========
==========
==========
==========
==========>] 610.95
  627KB/s    in 1.0s

2018-11-15 10:49:38 (627 KB/s) - 'c.txt' saved [625611/625611]

remasis@michelangelo:~/tmp/scinet/malware$ md5sum *
1765193e92c10eab1dc09a2c89857734  c.txt
ff0b04fcabdf6f14e324c1a9708d2a5  d.txt
remasis@michelangelo:~/tmp/scinet/malware$ file *
c.txt: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), statically linked, for
d.txt: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), statically linked, for
remasis@michelangelo:~/tmp/scinet/malware$
```

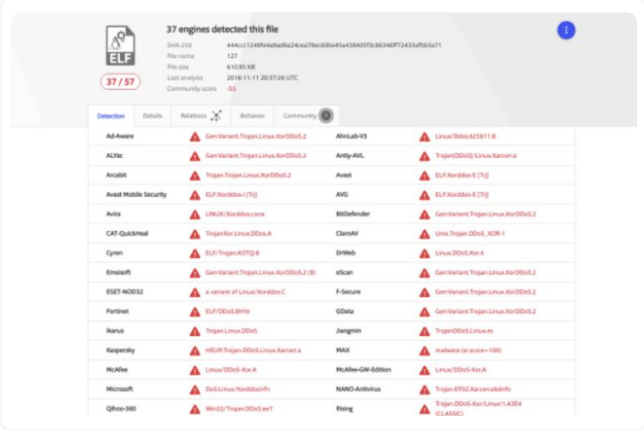

So, now that we know we have a compromised system and some idea about C&C servers, let's make sure no other IPs have been reaching out to these systems. We'll run a search over the past week of the conference to look for such activity:

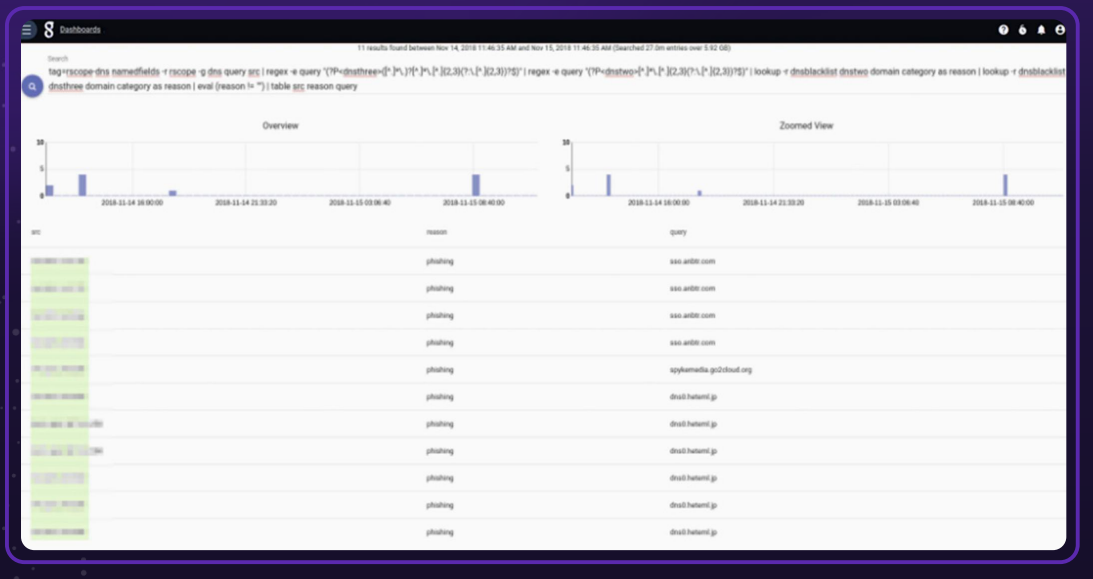

Thankfully, the only two hosts are the ones we have already identified from the tip .

We still don't know how they got compromised and we aren't 100% sure that they haven't conducted lateral movement, but at least we aren't seeing C&C traffic from any other systems . We can relax a little bit .

STOPPING ATTACKER EGRESS

Traffic blocking rules were put in place to prevent attackers from continuing the traffic egress. Basic traffic monitoring charts confirm correct application or rules and discontinued egress traffic

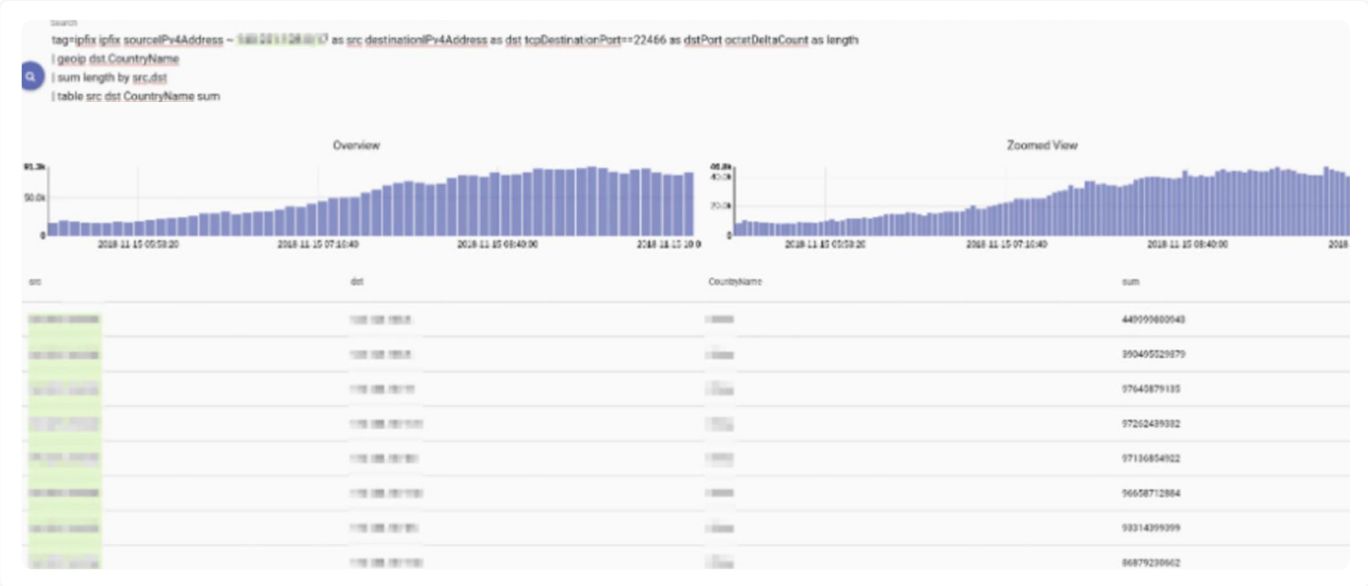

LATERAL MOVEMENT CONFIRMATION

No apparent lateral movement (connections from compromised machines to other machines in SC address space). The following query was used for both IPs over the last 2 weeks:

```
tag=rscope-conn namedfields -r rscope -g Conn conn_state src== “xxx.xxx.xxx.xx1” dst src_port
dst_port
| ip dst ~ xxx.xxx.xxx.0/17
| lookup -r iplist src address network as srcnet
| lookup -r iplist dst address network as dstnet
| table src dst dst_port srcnet dstnet
```

ROOT CAUSE ANALYSIS

To figure out the initial infection point, we conducted forensics both on device and in network data.

Host Based Forensics

'lsof -i +M' shows a mysterious process running.

Conveniently, that process does not show up in "ps auxwww" output. Also, it doesn't show up in a 'find' search for the filename, but we know the process id so all is not lost, let's look at /proc directly...

```
ls -i +M
COMMAND PID USER FD TYPE DEVICE SIZE/OFF NODE NAME
 sdf3fslsd 474 root 3u IPv4 512988 0t0 TCP hostname.redacted.sc18.org:39766- >ipxxx.ip-xxx-xxx-
xxx.yy:1522 (ESTABLISHED)
"sdf3fslsd"
```

Bingo. We've copied that binary off to another host for later investigation.

Now, hopefully they haven't re-written the logs and we can figure out when/ how the compromise happened

```
root@hostname.redacted.sc18.org:/proc/474#
ls -la total 0
 -r--r--r--
 lrwxrwxrwx 1 root root 0 Nov 15 12:34 cwd -
> /
 -r-------- 1 root root 0 Nov 15 12:34
environ
lrwxrwxrwx 1 root root 0 Nov 15 12:29 exe ->
/bin/sdf3fslsdf13 dr-x------ 2 root root 0
Nov 15 12:34 fd
dr-x------ 2 root root 0 Nov 15 12:34 fdinfo
```

gravwell

REMEDATION

The initial foothold was gained through brute force attacks against the SSH service on the system.

In this case, the ACL for SSH access needed to be strengthened to ensure this system could not be reached from the outside. This was already remediated at time of investigation, so no further action was required and the hunt was concluded.

CONCLUSION

The event was incredible and the entire Gravwell team had a blast working with some fantastic people. The nature of the event was beneficial for us for a few reasons.

First, the academic and public nature means we can create materials like this to serve as references for what is possible with Gravwell. This case study serves as a shining example of what proactive threat hunting can do in terms of detecting threats and reducing response time. With Gravwell, the SCinet Network Security team was able to detect and respond to a real attack in a matter of minutes instead of the 206 days that is average for US companies1.

Second, the high-performance computing environment, while not very similar to the average corporate infrastructure, does pose scalability challenges not seen by many of even the largest organizations. This gave us an opportunity to really flex the capabilities that we've built over the past years of development and demonstrate to the community what analytics engineered for modern computing can do.

The high-intensity event shook out some usability bugs for sure, but the infrastructure never faltered and Gravwell was able to provide exceptional analytics capabilities used by the whole team. It was a big win for us and we were very thankful for the opportunity. Huge thanks to the entire team and all of the volunteers who worked with us to make the conference a smashing success.

https://newsroom.ibm.com/2018-07-11-IBM-Study-Hidden-Costs-of-Data-Breaches-Increase-Expenses-for-Businesses

gravwell

Key takeaways

- For the Conference, Gravwell provided our analytics platform to the Network Security team
- Responsible for cyber security on a network consisting of $52 million in contributed hardware, software, and services plus 4 .02 Terabits per second of external capacity
- The network was made possible by the contributions of 40 industry- leading organizations, who in total donated $52 million in hardware, software, and services
- Gravwell ingested over 4 .6 billion entries comprising over 1TB of data from a variety of sources
- Analysts ran 4281 manual searches, 17325 automated searches, and viewed dashboards 1159 times during the two weeks in the Network Operations Center (NOC)
- The Network Security team used Gravwell to stop continuous internet attacks automatically — freeing up time to better to identify, hunt, and respond to an actual attack that sought to bring the entire force of 4 .02 Tb/s against an unsuspecting SaaS company
- With Gravwell, the Network Security team was able to detect and respond to a real attack in a matter of minutes instead of the 206 days that is average for US companies

CONTRIBUTORS AND SPECIAL THANKS

Thanks to Michael “Dop” Dopheide and Scott Chevalier for helping on the hunt outlined here.

Special thanks to the SCinet Network Security team.

Thanks and well done to the entire SCinet team who kept things operational at blazing speeds.