



gravwell

Gravwell in the SCinet NOC

AN AFTER ACTION REPORT
AND CASE STUDY IN THREAT HUNTING

SCinet





TABLE OF CONTENTS

Executive Summary	3
Background	4
Architecture	5
SCinet Network Security	7
The Job	8
Threats and alerts	9
Data Sources.....	9
Data Enrichments	10
Analysis Techniques.....	10
Example Autonomous SOC/NOC Activity	11
Findings.....	12
Case Study: Hunting Adversaries with Gravwell.....	14
The Tip.....	15
The Hunt	16
Stopping attacker egress	21
Lateral Movement confirmation	21
Root cause analysis	22
Host Based Forensics	22
Network Based Forensics	23
Malware Analysis	25
Remediation	26
Conclusion.....	26
Key Takeaways	27
Contributors and Special Thanks	27



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 SCinet?

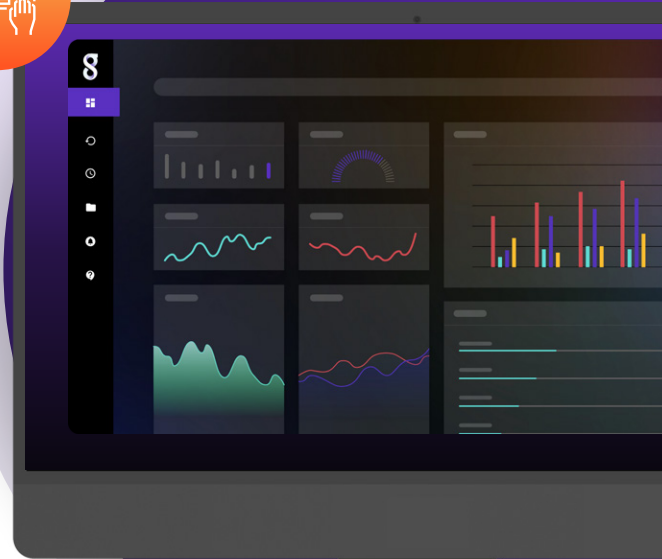
SCinet is the SC Conference's dedicated high-capacity network infrastructure, designed and built by volunteer experts from industry, academia, and government.

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 the 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:

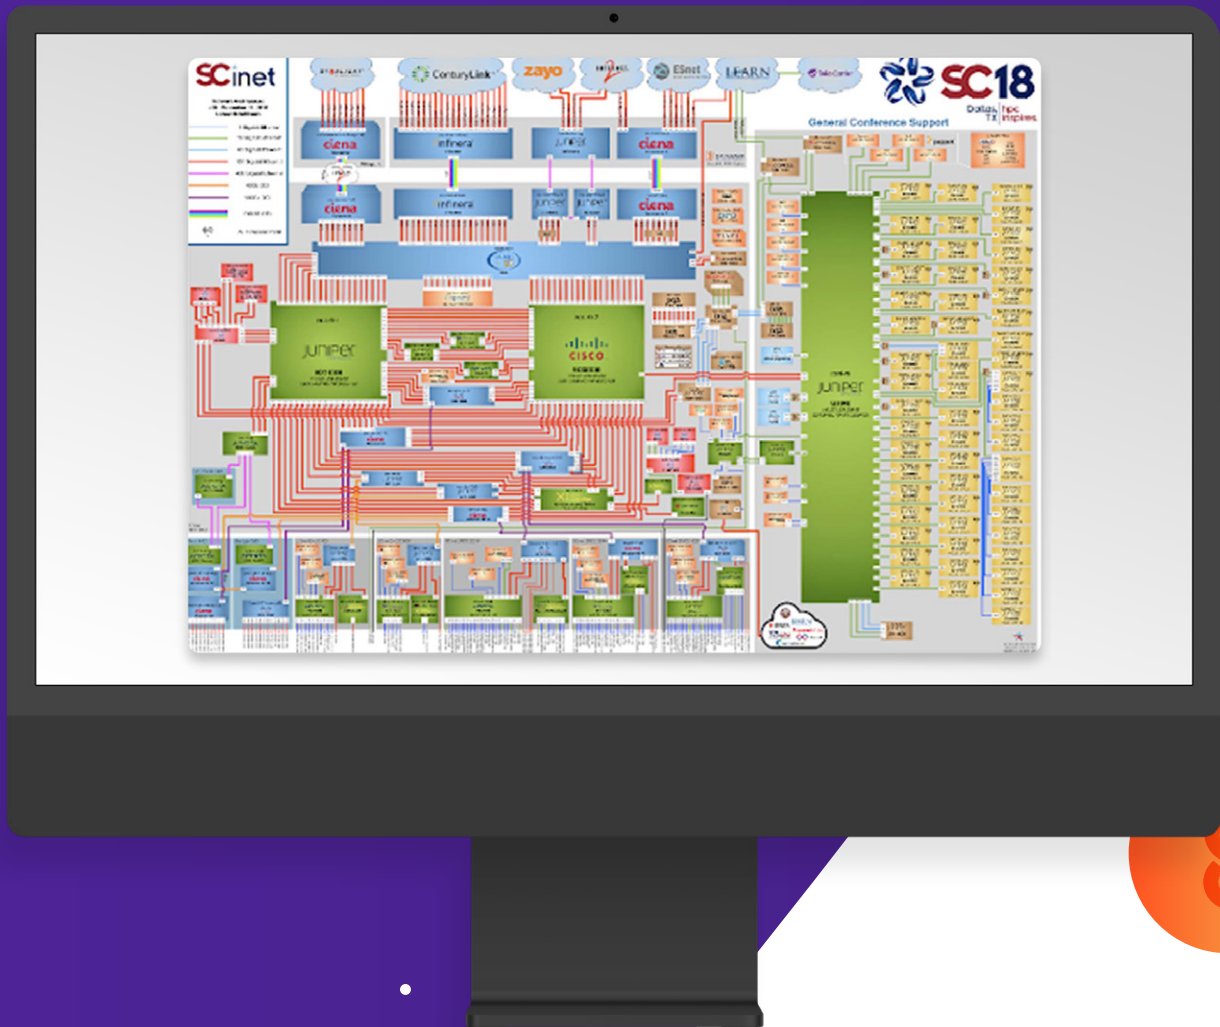
[www.youtube.com/
watch?v=B26DCSCL-7Q](https://www.youtube.com/watch?v=B26DCSCL-7Q)



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.





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 SCinet Network Security team breaks down the threat and risk landscape for the event to keep focus on providing the best possible experience for attendees, vendors, and researchers. 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 ENRICHMENTS

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.

If this sounds interesting, you might want to check out

<https://www.gravwell.io/blog/auditing-dns-with-coredns-and-gravwell>.

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.



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.

We configured the Gravwell automated scripting system to trigger on Attivo alerts and POST to an API on a bhr device (<https://github.com/ncsa/bhr-site>) which was the central authority for blacklists to block attackers.

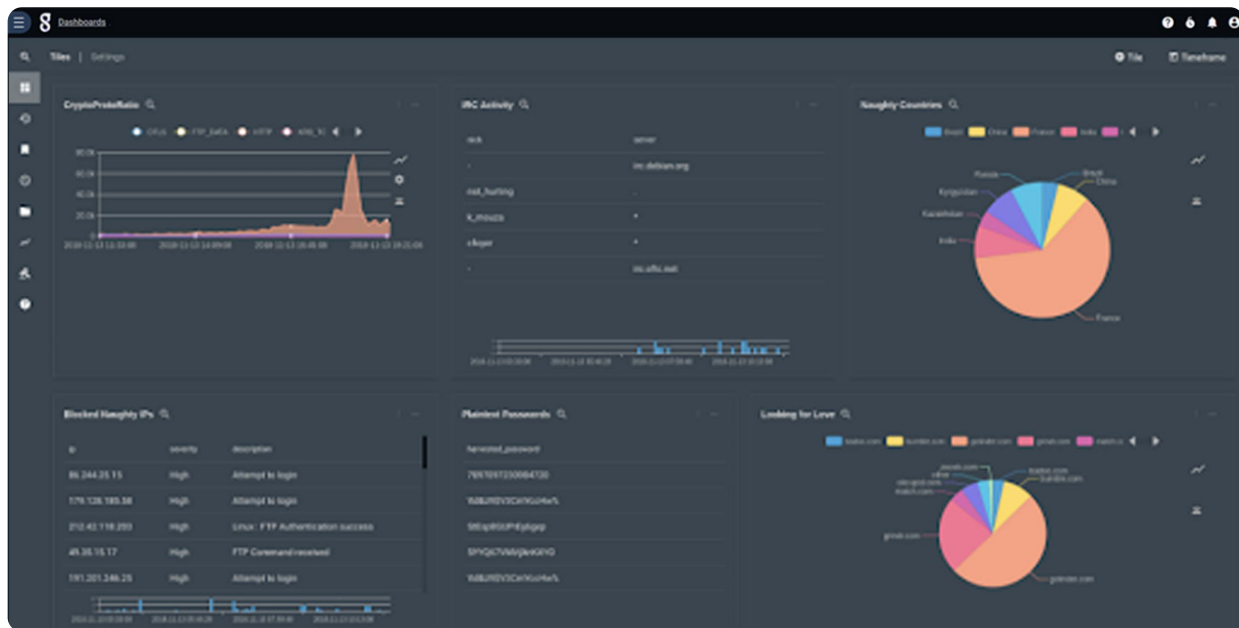
Thus, utilizing a combination of technologies, we could customize automation of tedious SCinet Network Security work to our needs and free up valuable analyst time to work on more dynamic and challenging problems. This one example of automating Gravwell + Attivo resulting in blocking hundreds of IPs 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.



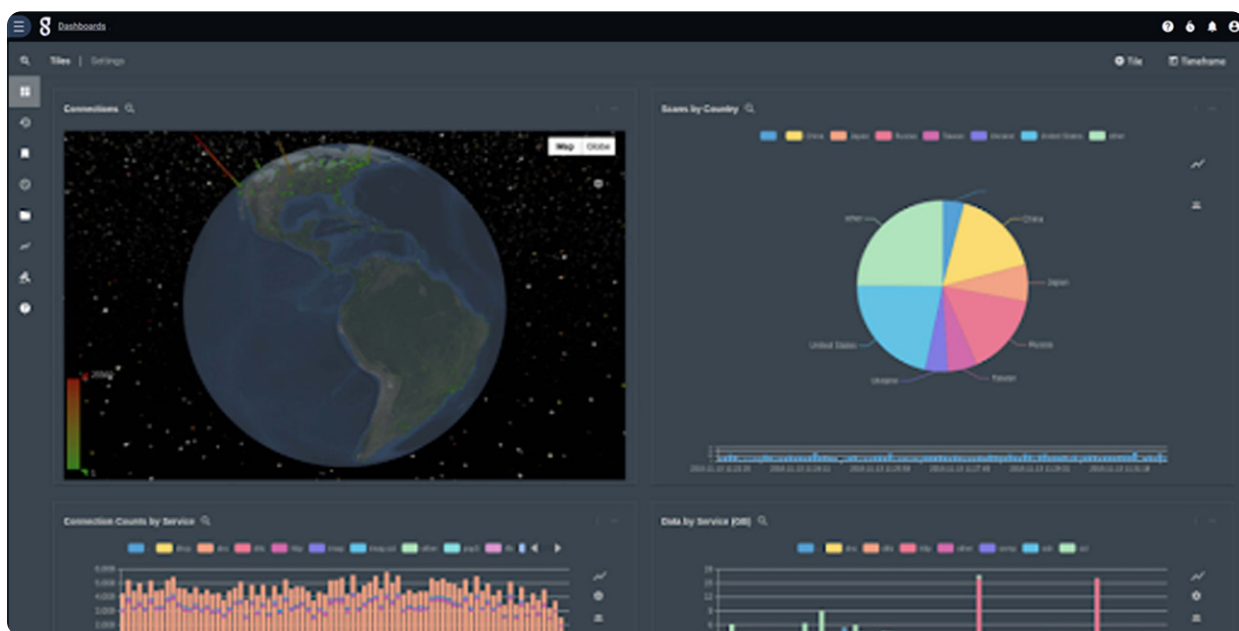
Findings

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.

Sidenote: These dashboards were created using the upcoming Gravwell user interface improvements set to drop in January 2019.



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.

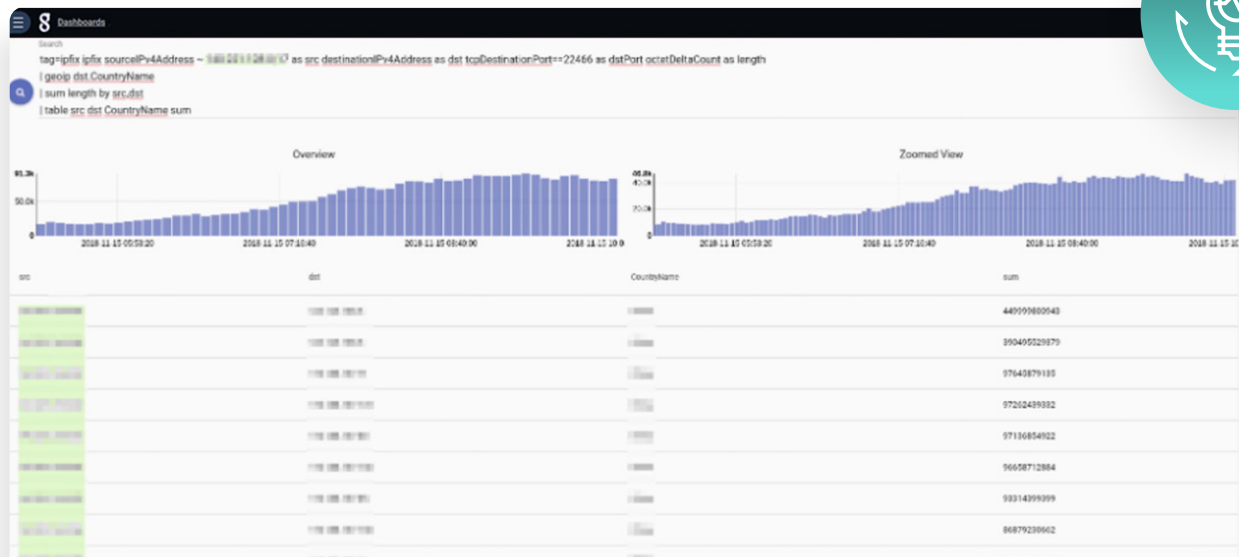
The bulk of the investigation was conducted utilizing Bro logs generated by Reservoir Labs' R-Scope. We also included data sources such as IPFIX, system logs, and data enrichment like a DNS blacklist from <http://www.malwaredomains.com/>.

Note: 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:



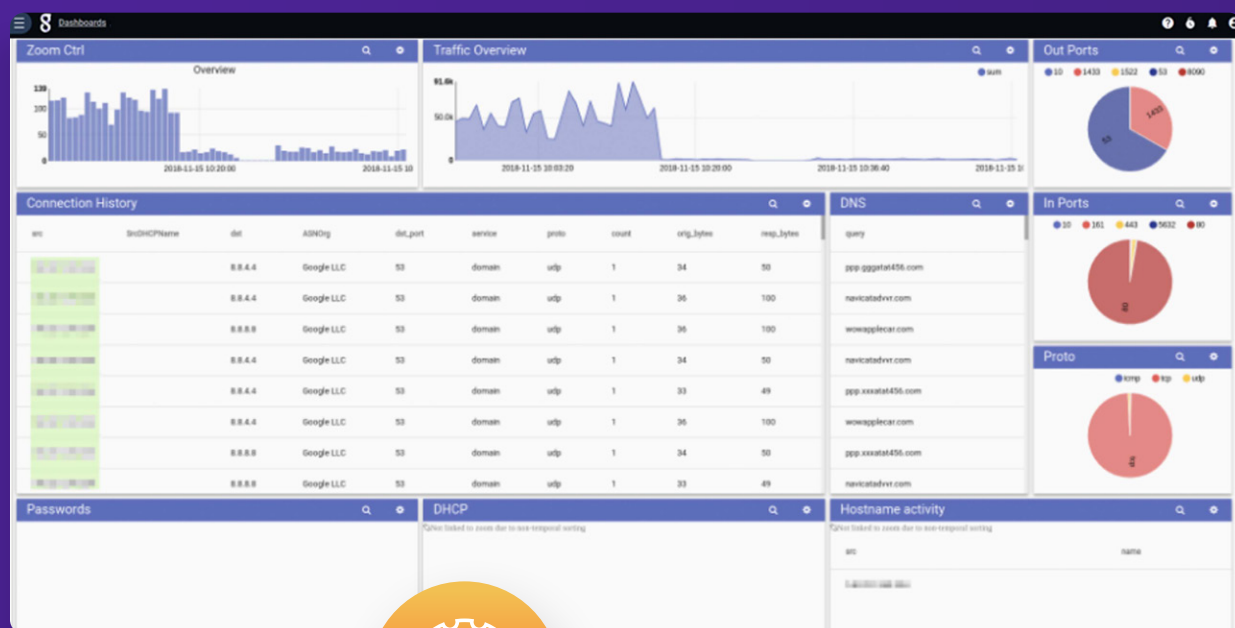
tag=ipfix ipfix sourceIPv4Address ~ xxx.xxx.xxx.0/17 as src destinationIPv4Address as dst
tcpDestinationPort==22466 as dstPort octetDeltaCount as length

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

11 results found between Nov 14, 2018 11:46:35 AM and Nov 15, 2018 11:46:35 AM (searched 27.0m entries over 5.92 GB)

Search

```

tag ← rscope dns namedfields ← rscope -g dns query src | regex -e query "(?P<dnstree>[^\.\?]{2,3}){1,3}|(?P<dnstree>[^\.\?]{2,3}){1,3}" | regex -e query "(?P<dnstree>[^\.\?]{2,3}){1,3}|(?P<dnstree>[^\.\?]{2,3}){1,3}" | lookup -r dnstree domain category as reason | eval (reason != "") | table src reason query
  
```

Overview

Zoomed View

src	reason	query
192.168.1.1	phishing	sso.ambit.com
192.168.1.2	phishing	sso.ambit.com
192.168.1.3	phishing	sso.ambit.com
192.168.1.4	phishing	sso.ambit.com
192.168.1.5	phishing	spykmedia.gocloud.org
192.168.1.6	phishing	drat.hetemi.jp
192.168.1.7	phishing	drat.hetemi.jp
192.168.1.8	phishing	drat.hetemi.jp
192.168.1.9	phishing	drat.hetemi.jp
192.168.1.10	phishing	drat.hetemi.jp

224 results found between Nov 13, 2018 11:30:07 PM and Nov 15, 2018 10:30:07 AM (searched 7.1m entries over 2.56 GB)

Search: `tag=scope:http namedfields + rscope -g HTTP src=*180 1 2% 25% host uri method status_code | table host method status_code uri`

The figure displays two bar charts. The 'Overview' chart shows the distribution of search results across time, with a peak around 2018-11-14 04:00:00. The 'Zoomed View' chart provides a more detailed look at the data, showing a significant peak around 2018-11-14 04:00:00 and another peak around 2018-11-14 10:00:00.

host	method	status_code	uri
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
www1.gggstat556.com	GET	200	/dd.nr
www1.gggstat556.com	GET	200	/dd.nr
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records
35.237.255.214	POST	200	/lookup/records

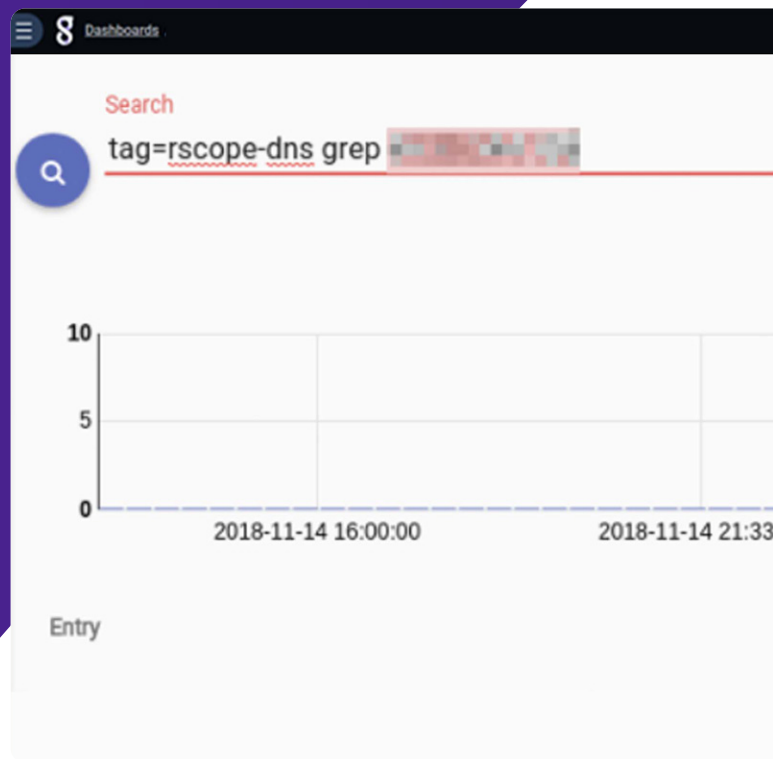
I reached out to grab the payload from the server using wget:

```
remasis@michelangelo:~/tmp/scinet/malware$ md5sum dd.rar
5fc38cdccd16710969b09582c9af34dc dd.rar
remasis@michelangelo:~/tmp/scinet/malware$ file dd.rar
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..5/1..GB..
00000070: 7145                                     qE
remasis@michelangelo:~/tmp/scinet/malware$
```

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. 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:

```
remasis@michelangelo:~/tmp/scinet/malware$ wget http://10.10.10.10/c.txt
--2018-11-15 10:49:37-- http://10.10.10.10/c.txt
Connecting to 10.10.10.10:80... 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 *
3765193e92c10eab1dc09a2c89857734 c.txt
7ff0b04fcabdf6f14e324c1a9708d2a5 d.txt
remasis@michelangelo:~/tmp/scinet/malware$ file *
c.txt: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), statically linked, for GNU/Linux 2.6.9, not stripped
d.txt: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), statically linked, for GNU/Linux 2.6.9, not stripped
remasis@michelangelo:~/tmp/scinet/malware$
```

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:

37 engines detected this file			
SHA-256	444c021249f649a8a24c0a79ec0be45a438405f3d86340f72433af8b3a71		
File name	127		
File size	610.95 KB		
Last analysis	2018-11-11 20:37:26 UTC		
Community score	-55		
Detection	Details	Relations	Behavior
Ad-Aware	GenVariant.Trojan.Linux.KorDDoS.2	AhriLab-V3	Linux/DoS.623611.B
ALYac	GenVariant.Trojan.Linux.KorDDoS.2	Antiy-AVL	Trojan(DDoS)/Linux.Xarcana.a
Arcabit	Trojan.Trojan.Linux.KorDDoS.2	Avast	ELF/XorDDoS-E [Trj]
Avast Mobile Security	ELF/XorDDoS-E [Trj]	AVG	ELF/XorDDoS-E [Trj]
Avira	LINUX/XorDDoS.coma	BitDefender	GenVariant.Trojan.Linux.KorDDoS.2
CAT-QuickHeal	TrojanXor.Linux.DDoS.A	ClamAV	Linux.Trojan.DDoS_XOR-1
Cyren	ELF/Trojan.KOTQ-6	DrWeb	Linux.DDoS.Xor.4
Emnisoft	GenVariant.Trojan.Linux.KorDDoS.2 (B)	eScan	GenVariant.Trojan.Linux.KorDDoS.2
ESET-NOD32	a variant of Linux/XorDDoS.C	F-Secure	GenVariant.Trojan.Linux.KorDDoS.2
Fortinet	ELF/DDoS.BHtr	GData	GenVariant.Trojan.Linux.KorDDoS.2
Ikarus	Trojan.Linux.DDoS	Jiangmin	Trojan.DDoS.Linux.es
Kaspersky	HEUR/Trojan.DDoS.Linux.Xarcana.a	MAX	malware (ai score=100)
McAfee	Linux/DDoS-Xor.A	McAfee-GW-Edition	Linux/DDoS-Xor.A
Microsoft	DoS.Linux/XorDDoS.cofn	NANO-Antivirus	Trojan.EF32.Xarcana.ebdfnfo
Qhoo-360	Win32/Trojan.DDoS.eef	Rising	Trojan.DDoS-Xor.Linux/1.A3E4 (CLASSIC)

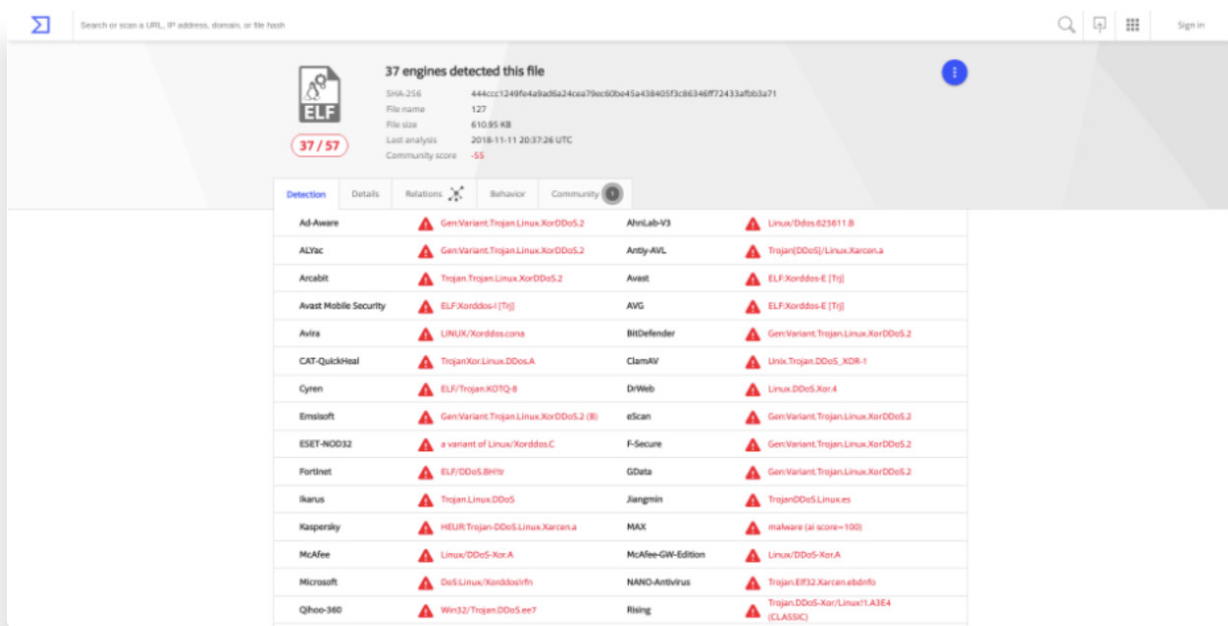
So, now that we know we have a compromised system and some idea about C&C servers, let's make sure no other IPs have been reaching out to these systems. We'll run a search over the past week of the conference to look for such activity:



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 of 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

'ls -i +M' shows a mysterious process running:

```
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"
```

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...

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
dr-xr-xr-x  9 root root 0 Nov 15 12:29 .
dr-xr-xr-x 181 root root 0 Nov 15 12:29 ..
dr-xr-xr-x  2 root root 0 Nov 15 12:34 attr
-rw-r--r--  1 root root 0 Nov 15 12:34 autogroup
-r-----  1 root root 0 Nov 15 12:34 auxv
-r--r--r--  1 root root 0 Nov 15 12:34 cgroup
--w-----  1 root root 0 Nov 15 12:34 clear_refs
-r--r--r--  1 root root 0 Nov 15 12:30 cmdline
-rw-r--r--  1 root root 0 Nov 15 12:34 comm
-rw-r--r--  1 root root 0 Nov 15 12:34 coredump_filter
-r--r--r--  1 root root 0 Nov 15 12:34 cpuset
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
```

The current auth.log doesn't show anything useful, but looking at bit back in time we have a winner:

```
root@hostname.redacted.sc18.org:/var/log# grep ssh auth.log.1 |grep Accepted
Nov  4 18:55:12 hostname.redacted.sc18.org sshd[24387]: Accepted password for root
from xxx.xxx.xxx.xxx port 53428 ssh2
root@hostname.redacted.sc18.org:/var/log# zcat auth.log.2.gz |grep ssh |grep Accepted
Oct 29 19:01: hostname.redacted.sc18.org sshd[17348]: Accepted password for root
from xxx.xxx.xxx.xxx port 58190 ssh2
Oct 29 20:45:01 hostname.redacted.sc18.org sshd[18879]: Accepted password for root
from yyy.yyy.xxx.xxx port 42372 ssh2
Oct 29 20:47:53 hostname.redacted.sc18.org sshd[19020]: Accepted password for root
from zzz.zzz.zzz.zzz port 39255 ssh2
Nov  2 19:04:52 hostname.redacted.sc18.org sshd[29293]: Accepted password for root
from xxx.xxx.xxx.xxx port 46072 ssh2
```

The attacker IP of **zzz.zzz.zzz.zzz** originated from an overseas country to which the vendor has no affiliations.

We believe the vendor VM was brought up with a default password and connected to the internet prior to SCinet "go live". Since then, it's been constantly downloading dd.rar. The root password was changed by the vendor upon SCinet "go live" so the SCinet Network Security ssh-auditor scanning didn't catch the vulnerability. Thus, the initial exploitation flew under the radar.

NETWORK BASED FORENSICS

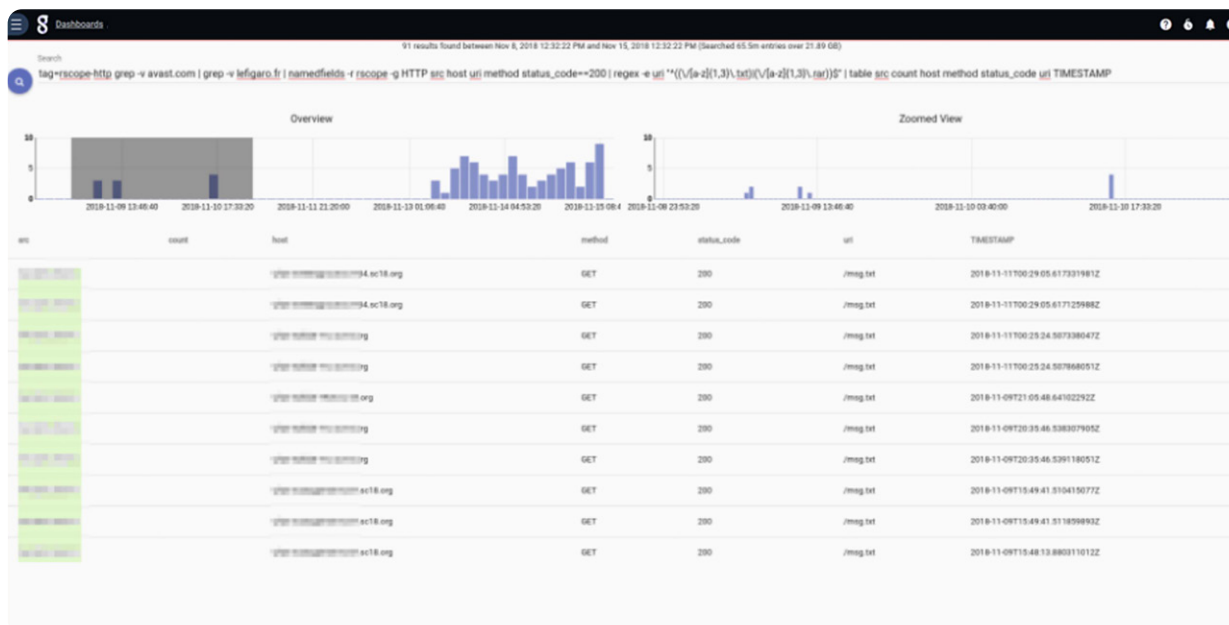
Based on the observed patterns we can develop and Indicator of Compromise (IOC) or a reachout to the URL pattern of one to three letters followed by .txt or .rar. I.e. matching the following regex:

```
"^((\/[a-z]{1,3}\.txt)|(\/[a-z]{1,3}\.rar))$"
```



This does result in a few innocuous entries so we will add some manual filtering at the front of the query with grep -v.

The resulting query looks something like:

```
tag=rscope-http grep -v avast.com | grep -v lefigaro.fr | namedfields -r rscope -g HTTP
src host uri method status_code==200 | regex -e uri "^((\/[a-z]{1,3}\.txt)|(\/[a-z]{1,3}\.
rar))$" | table src count host method status_code uri TIMESTAMP
```

Using the zoom feature on the entries we can see that the early activity is innocuous. Starting at around 2018-11-13T18:42:18.523044109Z is when the compromise is clearly visible.



Just to be sure, let's clean up the query a bit and look even longer. If we explicitly only search for an HTTP GET on /dd.rar, the results are:

Uh oh. That seems bad. There are other IP addresses a full 2 weeks earlier. [xxx.xxx.xxx.xx1](#) and [xxx.xxx.xxx.xx2](#). Based on the host-based forensics and lack of lateral movement, however, it would be reasonable to suggest that these are the same machines and that they had an IP change.

MALWARE ANALYSIS

The persistence mechanism employed by the malware was an insertion into cron to run hourly.

The malware executes the following shell command:

```
sed -i '/\/etc\/cron.hourly\/gcc.sh/d' /etc/crontab && echo '*/*/*/* root /etc/cron.hourly/gcc.sh'
```

The persistence appears to be successful as we can see network effects from the execution on even 1-hour intervals.

src	uri	TIMESTAMP
	/dd.rar	2018-11-15T17:17:37.360296964Z
	/dd.rar	2018-11-15T16:17:31.630299091Z
	/d.txt	2018-11-15T15:39:39.064129114Z
	/dd.rar	2018-11-15T15:17:25.523960113Z
	/dd.rar	2018-11-15T14:17:19.482155084Z
	/dd.rar	2018-11-15T13:17:14.122636079Z
	/dd.rar	

Unfortunately, we left our IDA at home so we're doing this old school with objdump.

Call to install the crontab entry.

```
804d08a: 00 00 00          jmp 804dfdd <main+0x1459>
804d08d: e9 4b 0f 00 00    mov -0xc(%ebp),%eax
804d092: 8b 45 f4          movl $0x0,(%eax)
804d095: c7 00 00 00 00 00 movl $0x0,(%eax)
804d09b: 8b 45 f4          mov -0xc(%ebp),%eax
804d09e: 89 04 24          mov %eax,(%esp)
804d0a1: e8 4a dd 01 00    call 806adf0 <shmdt>
804d0a6: 8b 8d 40 c2 ff ff mov -0x3dc0(%ebp),%ecx
804d0ac: 8b 41 04          mov 0x4(%ecx),%eax
804d0af: 8b 00             mov (%eax),%eax
804d0b1: 89 04 24          mov %eax,(%esp)
804d0b4: e8 34 be ff ff    call 8048eed <LinuxExec>
804d0b9: 8b 45 d0          mov -0x30(%ebp),%eax
804d0bc: 89 04 24          mov %eax,(%esp)
804d0bf: e8 2b bd ff ff    call 8048def <DelService_form_pid>
804d0c4: c7 85 44 c2 ff 00 movl $0x0,-0x3dbc(%ebp)
804d0cb: 00 00 00          jmp 804d0c4
```

REMEDIATION

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.

¹ <https://newsroom.ibm.com/2018-07-11-IBM-Study-Hidden-Costs-of-Data-Breaches-Increase-Expenses-for-Businesses>

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.