

Copyright!

(U) utility of "security conferences"



(hide)

2012.02.08 - 09:53 am

**Do not reuse
this image!**

(U//FOUO) Caveat: it's been a couple years since I have attended a security conference(Shmoocon/Defcon/Blackhat). My opinions are not formed from recent first-hand experience, but through previous stimuli to my cerebrum that have been confirmed by 2nd hand experience over the last couple of years. When I first went to Blackhat/Defcon, it was with the wide-eyed anticipation of, "I'm going to go listen to all of the talks that I can, soak up all of the information possible, and become a supar-1337-haxxor." What a let-down of an experience that was. You find the most interesting topics and briefings, wait in lines to get a seat, and find yourself straining your ears to listen to someone that has basically nothing new to say. Most of the talks get hyped up exponentially past any amount of substance they actually provide, most of the "interactive sessions" end up in a "oh! woe is the state of the security industry!" chant, and leave the audience no better off than before.

(U//FOUO) If you want to learn crazy new things, more often than not, you won't find it at a talk in a con. Why not walk around NSA, find people in offices that do things you find interesting, and talk to them about how they do what they do(or find a mentor in that area)? Despite stereotypes of the kinds of people that work here, many people are kind and open enough to share their trade-craft with others. We have the luxury of working in a community that has some of the brightest, smartest, and most cutting edge people around, it would be a shame for people to constantly attend cons hoping to learn that "cool new thing", when there is exponentially more knowledge sitting around them every single day at work.

(U//FOUO) Granted, there are always a couple exceptional talks at the cons, but, in my humble opinion, they don't make up for the overall lack of content. So, what good are these conferences? My personal opinion is that their utility is mainly for social interaction and meeting *relatively* like-minded individuals. It's the ability to kick back for a weekend and geek-out with other people. For some, this makes the cost of the con completely worth it, others may be severely disappointed...it all depends on what you expect to get out of it.

Current Mood: 😊 okay[Leave a comment](#)

Copyright!

(U) I hunt SIGINTs (part 1)



(hide)

2012.12.12 - 10:10 am

**Do not reuse
this image!**

(S//SI//REL) So, SIGINT is down right cool. As much as we complain about our "Big Data Problem", collection/processing issues, dismal infrastructure/outdated browsers/OS's, our ability to pull bits out of random places of the Internet, bring them back to the mother-base to evaluate and build intelligence off of is just plain awesome!

(S//SI//REL) One of the coolest things about it is **how much** data we have at our fingertips. If we *only* collected the data we knew we wanted...yeah, we'd fill some of our requirements, but this is a whole world of possibilities we'd be missing! It would be like going on a road-trip, but wearing a blindfold the entire time, and only removing it when you're at one of your destinations...yeah, you'll still see stuff, but you'll be missing out on the entire journey!

(S//SI//REL) So I decided to write a short series (affectionately titled '*I hunt ...*') on things that I'm trying to do with data that wouldn't normally be interesting by itself, but by thinking about it in a new way, makes it extremely valuable. My interests lately have been in using passive collect to identify/enable CNE efforts, so that's predominantly what the first few topics will be about.

(U//FOUO) If there are any topics someone wants to see specifically, let me know. As well, if any of the following information is useful, please let me know and I can put more out. **Part 2 - Hunting sys admins** coming very soon!

2 comments :: [Leave a comment](#)

Copyright!

(U) I hunt sys admins (part 2)



 (hide)

2012.12.12 - 12:42 pm

Do not reuse
this image!

Entry tags: admin, cne, infrastructure, quantum, target

(U//FOUO) *This post is meant to provide a background for ***why*** it's good to target sys admins in SIGINT. If you already know this, feel free to skip forward to the next sections.*

(S//SI//REL) Being in SID, our overall goal is to produce intelligence to give to decision-makers. How we go about doing that, is whenever a target uses technology to communicate, we collect it, analyze it, and write reports on it. Sounds simple enough...except for the fact that we have to be targeted in what networks we collect. We can't collect everything all the time, so if a target starts to communicate on a network where we are **not** collecting, there is some manual leg-work that has to be done to steer the SIGINT system in their direction. This is where I must introduce my loyal friend, the sys admin.

(S//SI//REL) Up front, sys admins generally are not my end target. My end target is the extremist/terrorist or government official that happens to be using the network some admin takes care of. Sys admins are a means to an end. For example, assume your target is using a CDMA device on a foreign network; there may be situations where we passively collect his phone call/SMS out in the wild, but it would be ***really*** nice if we had access to the local infrastructure where we could monitor which tower he's connected to at any given point in time, or monitor all phone calls/data traffic that his phone generates. Many times, it's difficult to directly target infrastructure...generally we'll need a fair amount of information going into an operation, such as:

- * topology of the network we are targeting
- * credentials for infrastructure devices
- * situational knowledge, such as access lists set up to only allow specific IP addresses to administer certain machines
- * an overall knowledge of how the network is put together and configured

(S//SI//REL) In order to get that, who better to target than the person that already has the 'keys to the kingdom'? Many times, as soon as I see a target show up on a new network, one of my first goals is, "Can we get CNE access to the admins on that network, in order to get access to the infrastructure that target is using?"

"Yeah, that pretty much makes sense, but how are you 'just gonna get CNE access' on an admin?"

(S//SI//REL) Good question, thanks for asking. Most of the time I'm going to rely on [QUANTUM](#) to get access to their account (*yeah, you could try spam, but people have been getting smarter over the last 5-10 years...it's not as reliable anymore*). So, in order to work our QUANTUM-magic on an admin, we'll need some sort of webmail/facebook selector for them.

*"You know, you ***could*** just look up the 'point of contact' in the registry information associated with their IP space/domain names..."*

(S//SI//REL) Yeah, you could do that. Personally, I haven't had a huge amount of luck with it, because most of the time I end up running across their ***official*** e-mail address that's hosted on their own network. That's generally not a recipe for success in the QUANTUM world, what we'd **really** like is a personal webmail or facebook account to target. There's a couple ways you could try this: dumpster-dive for alternate selectors in the big SIGINT trash can, or pull out your wicked Google-fu to see if they've posted on any forums and list both their official and non-official e-mails in a signature block...but what if there was another way to do it?

(S//SI//REL) Other fun (read:useful) things to get off of a sys admin (from my point of view):

- * network maps off of their hard drive
- * credentials from text files (or from our key-loggers...potato potato)
- * full lists of customers (along with associated dedicated IP allocations is a bonus)
- * e-mail with upstream providers detailing how your network is connected to the bigger Internetz. For example, if I see they use certain fiber cables to connect to the world, I'll go look in [SSO](#)'s collect for their traffic. If they use VSAT's, I'll go look for their network in [FORNSAT](#)'s environment.
- * pictures of cats in funny poses with amusing captions

(S//SI//REL) But **all** of this boils down to getting an admin's webmail/facebook account in order to QUANTUM it and get CNE access to their box. Next section will detail targeting admins who use telnet...

[3 comments](#) :: [Leave a comment](#)

Copyright!

(U) I hunt admins that use telnet (part 3)




 (hide)

2012.12.12 - 04:09 pm

Do not reuse this image!

Entry tags: access list, admin, cloud, cne, passive, quantum, sigint, telnet

(S//SI//REL) If a target that I care about is on a network that I don't have access to, in [this](#) post I described that I will try to get access to that network by targeting the sys admin. In order to target the sys admin, it's easiest if I know what their personal webmail/facebook username is so that I can target it with [QUANTUM](#). The hardest part is identifying that admin's personal account to target in the first place.

Now, fade off with me into dream-land. Pretend that we had some master list. This master list contained tons of networks around the world, and the personal accounts of admins for each of those networks. And any time you wanted to target a new network, you could just find the admin associated with it, queue his accounts up for QUANTUM, get access to his box and proceed to pwn the network. Wouldn't that be swell?

(S//SI//REL) Well, you can stop dreaming my friends, I think it's possible (at least kinda partially). And we'll get started on this endeavor by chasing down admins that use telnet. By this point, I'm hoping you're saying, "Telnet? Telnet?! No one should still be using telnet!" That is the correct response, however, telnet (as an administrative tool on the Internet) is alive and well. In fact, it's so alive and well, [DISCOROUTE](#) is a tool specifically designed to suck up and database router configuration files seen in passively collected telnet sessions (for the record, DISCOROUTE is awesome, and you should check it out if you have at least one iota of love for SIGINT). So, what shall we do with all of these configuration files?

"Dude! Map all the networks!!!"

(S//SI//REL) Yes, that is definitely a useful thing to do with all of these router configs, and is something that will need to be done. But understanding the topology of a network isn't necessarily going to get us **access** to that network...per my previous post, I mentioned how getting CNE access to an admin is usually a golden ticket into the network. So, this raises a question:

"How can we use a router configuration file to find the personal webmail/facebook accounts of a sys admin?"

(S//SI//REL) To illustrate this point, I randomly picked a config for a router in Kenya...the link is NF, so send me an e-mail if you'd like to look at the config in it's entirety. For now, I'll just copy/paste the relevant portions. This process, as with any analysis at it's most fundamental level, is based on assumptions. Here are probably some safe assumptions:

- * We have a router, and that router has an admin
- * By the nature of having the config in DISCOROUTE, the admin uses telnet to log into the router
- * The admin probably doesn't want to let anyone and everyone else log into the router
- * The admin may set up an access control list to only allow his IP addresses to telnet to the router

(TS//SI//REL) Those seem like relatively safe assumptions to make for the moment. Now, let's see how those manifest in our example Kenyan config (if you're not familiar with router configs, don't worry, this will be pretty basic):

1) First, let's look at the vty (read:telnet) lines on the router, and see if there is an access list associated with them:

```
line vty 0 5
  access-class 11 in
  password 7 [REDACTED]
```

(TS//SI//REL) Yup, for all of the telnet lines, access list #11 is applied...which basically means, if you want to telnet to this router, you have to meet the criteria of that access list. (oh, before I forget, password 7's are ROFL-easy to crack. You can google 'cisco password 7 cracker' and get web pages that allow you to copy the password 7 hash, and it'll break it for free...anyone can figure out the password for this router. And pointing out for the lulz, the enable password for this router is password 7 hashed as well!)

2)...anyway, on with the access list #11:

```
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
access-list 11 permit [REDACTED]
!
```

(TS//SI//REL) Okay, this is relatively self-evident. If you want to telnet into this router, you have to be coming from one of those IP addresses that are permitted. Those who are slightly network-savvy will also be quick to point out, that even if you know the credentials for this router, and know which IP's are whitelisted, you **CAN NOT** just spoof the source IP of packets to try to log into this router...b/c you'd never see the responses. So you *actually* have to have access to one of these IP's to log into this router. Fair enough? Mmkay

(TS//SI//REL) I want to take a quick look back at our assumptions for a sec. We assumed that an admin would allow himself to telnet, but not others. So, based on the combination of that assumption with the info from the config, we can make the following assumption:

"The IP addresses in that access list probably belong to entities that administer that Kenyan network."

(TS//SI//REL) Hopefully that sounds fair enough. Here's the fun part...why don't we take those IP addresses, and look for *anyone* actively logging into their hotmail/yahoo/facebook/etc accounts from those IP addresses within the recent past? With whatever results you get back, you now have a **probable list of personal accounts of administrators for that Kenyan network!** From here, if you need CNE access to that network, just pull those selectors, queue them up for QUANTUM, and proceed with the pwnage. Yay! /throws confetti in the air

"Okay, that sounds relatively reasonable, but that's still a pretty manual process..."

(TS//SI//REL) Yeah, you're correct...but here's where "the cloud" actually comes to the rescue (yeah, I said it, crucify me)...All of the DISCOROUTE data is dumped into the GM Place cloud. All [ASDF](#) gets dumped into the GM Place cloud (ASDF is the metadata that gets generated for almost every session that we collect in our big bad passive SIGINT system). So, all someone would have to do is write a cloud analytic that would do the following:

- 1) parse through all of the router configs in DISCOROUTE
 - 1a) identify every router that has an access list tied to a vty (telnet) line
 - 1b) look through the access list, and pull out all public IP addresses
 - 1c) put those public IP addresses in a list somewhere
- 2) parse through all the Active User events in ASDF
 - 2a) take the list of public IP's that are associated with probable sys admins and use it as our seed list
 - 2b) look for all Active User login activity from those IP addresses
 - 2c) take all accounts that come back and put those into another list

(TS//SI//REL) So, by combining all of that information, you end up with a list of public IP addresses that probably belong to sys admins as well as personal accounts that probably belong to those admins. All you have to do is put all this info in a database somewhere, and what you end up with is a list of networks as well as personal accounts of probable admins for those networks! Then, as soon as one of those networks becomes a target, all TAO has to do is query the database, see if we have any admins pre-identified for that network, and if we do, automatically queue up tasking and go-go-CNE!

(TS//SI//REL) All of this can be done by tweaking the data that we **already have at our fingertips!!!** Remember, our "Big Data Problem" is that we have too much data...all we have to do is find ways of taking disparate data sets that wouldn't necessarily be interesting by themselves, but when you put them together in the right way become simply awesome!

"Yeah, okay, that's cool and all, but that relies on sys admins being bad and using telnet. *I* don't use telnet and most people that I know don't use telnet, it would never be able to identify us!"

(TS//SI//REL) Good point, this analytic does rely on admins using telnet. **HOWEVER**, I have an idea for how to identify personal accounts of admins that use SSH as well, which I'll talk about in my next post (I hunt admins that use SSH)!

Current Mood:scheming

[7 comments](#) :: [Leave a comment](#)

Copyright!

(U) I hunt admins that use SSH (part 4)



 (hide)

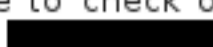
2012.12.13 - 02:41 pm

Do not reuse this image!

Entry tags: cne, quantum, router, server, ssh

(S//SI//REL) Welcome back comrade! It's good to see you again. For background, I talked in [this post](#) about why I target system administrators when I want CNE access to a network, and I talked in [this other post](#) about how we can get sys admin's personal selectors (for targeting) based on router configs that we passively collect through telnet sessions. This begs the obvious retort:

"Yeah, but that wouldn't work at all against any sys admin that uses SSH, because we would never see the content of the config! It's encrypted!"

(S//SI//REL) That is absolutely correct, however, [Options Exist!](#) I still think it's possible to identify a sys admin's personal account (sometimes) **even if** they use SSH. First off, if you're at all unfamiliar with SSH, please feel free to check out the resource graciously posted by one of my favorite high-side online personas,  located [here](#).

(S//SI//REL) Before we get into the methodology, let's talk about SSH for a minute, and some assumptions that we're going to make (up front caveat: I'm assuming we're looking for SSH on port 22. I'm sure there's ways to find it on non-standard ports, but just for the sake of clarity, I'll refer to SSH sessions as defined by being on port 22). Let's assume that you try to use SSH to connect to some server:

```
[madhatter@localhost ~]$ ssh 1.2.3.4
```

(U//FOUO) Assuming 1.2.3.4 actually has SSH listening on port 22, it'll go through the 3-way TCP handshake, and the server will send you back some packets and something like this may pop up:

```
Warning! This is a banner for some router owned by some company,
if you're not authorized to log in to this device, disconnect immediately!
Blah...blah...blah...
```

```
Enter username:
```

(U//FOUO) Here's where you enter your username...we'll pretend you type '**admin**' and press Enter. Then you get this:

```
Enter password:
```

(U//FOUO) Okay, here's where the rubber meets the road. One of two options exist. Either you have the right password, or you don't. Let's explore what would happen under each of the circumstances. First off, let's pretend you do **NOT** have the right password. You try '**admin**' for the password and press enter:

```
Invalid password. Enter password:
```

(U//FOUO) If you're password guessing, you'll probably get about 3 tries before the server kills the connection, and you have to reconnect and try again. However, if you **do** have the right password, you'll log in, you can run your commands to do whatever you want to do, and the server will spit back the output for all of the commands that you run.

(TS//SI//REL) I know, that sounds uber-simple, but here's why it matters. Think about what all of this would look like if we were to observe this in our passive SIGINT system. We would probably see encrypted traffic between two IPs, one of the ports would be on port 22 (we'll say that's the server side). The traffic going **TO PORT 22** (in the client-to-server direction) will consist of sending a username, sending a password, and sending commands (assuming the client successfully logged in). For traffic coming **FROM PORT 22**, well...that depends.

(S//SI//REL) If the client does **not** have the right credentials, we'd expect the server-to-client direction to consist of:

- 1) sending the device's banner
- 2) sending the prompt for a username
- 3) sending the prompt for a password
- 4) sending another prompt for a password
- 5) sending ANOTHER prompt for a password
- 6) subsequently giving up, killing the connection, and forcing the client to restart and try again

(S//SI//REL) If the client **does** have the right credentials, we might expect the server-to-client direction to consist of:

- 1) sending the device's banner
- 2) sending the prompt for a username
- 3) sending the prompt for a password
- 4 thru ?) sending back the output for whatever commands the client runs

(S//SI//REL) So, purely based off of the above assessment, we would expect **unsuccessful** logins to be consistently small (as in, the number of bytes) in the server-to-client direction. However, **successful** logins will be of variable length, but probably consistently larger in size (in bytes) when compared to unsuccessful logins. So, one assumption that I am suggesting is:

You can guesstimate whether an SSH session was successful or not *PURELY* based off of the size of the session in the server-to-client direction.

(S//SI//REL) So, imagine that you do some analysis, and determine that 1500 bytes is roughly a good number to differentiate between successful/unsuccessful SSH login attempts. Any server-to-client SSH session below that size is **probably** an unsuccessful attempt, and any server-to-client SSH session over that size is **probably** a successful attempt. What could you do, armed with information like that?

- 1) You could create lists of client IP addresses that **consistently** are unsuccessful in SSH sessions to multiple servers. Then you could put these in a, "probably password guessers/probably brute forcers" list.
- 2) You could also create lists of IP addresses that appear to be successful in having access to other IPs. (*ahhhh yeah, here's where we can have fun*)

(S//SI//REL) I'm sure there are a plethora of other things you can do with that sort of data, but I'm really interested in #2 at the moment. Based purely off of:

- * FROM port 22
- * session size is greater than 1500 bytes

then I can infer that:

- * To IP = admin
- * From IP = server/router
- * The admin appears to have successful access to the server

"Aaaaand?"

(TS//SI//REL) From here, all I have to do is recycle my methodology from last post! I can scour **ALL OF SIGINT** for sessions that meet the above criteria, harvest a list of "admin IPs", use that as a seed list to go back through all ASDF looking for any Active User login events within a time frame of the SSH session, and now I have personal accounts for people that are probably admins to the server IP. So if the server IP is ever in a network that I want access to, I don't have to decrypt the admin's SSH session, all I have to do is hope he checked his facebook/webmail within a certain timeframe of SSH'ing to the server. If he did, that selector is now tasked for QUANTUM, and we wait to get access to his box.

(S//SI//REL) One of the reasons why I'd only look for Active User logins within a certain timeframe of an SSH session, is that I want to make sure to find accounts that are actually associated with the admin, and not some other random person who happened to receive that same IP via DHCP the next day (or something like that).

(S//SI//REL) The previous few posts have gone into some fun ways that we can harvest the power of the SIGINTs to target sys admins of foreign networks. My next post will go into an uber-simple way that we can try to detect when people **other** than the legitimate sys admin have access to a router (*oh yeah, find me somma that cyber!*).

(S//SI//REL) I'm sure there are other fun and innovative ways to go about this, feel free to share them if you think of any. Also feel free to take any of the above thoughts/ideas and use them for your own purposes (like IAD maintaining a list of IP's that we see password guessing in SIGINT?). As always, if you have any questions, or want more info, feel free to drop me an e-mail or leave a comment!

Current Mood:juche-licious

21 comments :: [Leave a comment](#)

Copyright!

(S//SI//REL) I hunt people who hack routers (part 5)



(hide)

2012.12.14 - 01:08 pm

Do not reuse
this image!

Entry tags: cne, discoroute, router, sigdev

(TS//SI//REL) Happy Friday my esteemed and valued Intelligence Community colleagues! There has been a topic of conversation that has started to rumble beneath the surface of the Cyber-scene lately, it's about router hacking(for this post, I'm not talking about your home ADSL router, I'm talking about bigger routers, such as Ciscos/Junipers/Huaweis used by ISPs for their infrastructure). Hacking routers has been good business for us and our 5-eyes partners for some time now, but it is becoming more apparent that other nation states are honing their skillz and joining the scene. Before I get into it too much, let's go over some of the things that someone could do if they hack a router:

- * You could add credentials, allowing yourself to log in any time you choose
- * You could add/change routing rules
- * You could set up a packet capture capability...imagine running Wireshark on an ISP's infrastructure router...like a local listening post for any credentials being passed over the wire(!)
- * You could weaken any VPN encryption capabilities on the router, forcing it to create easily decryptable tunnels
- * You could install a dorked version of the Operating System with whatever functionality you want pre-built in

(TS//SI//REL) There are a plethora of things you could do once you get CNE access to a router...suffice it to say, getting access to a router is **very good** for the actor, and **very bad** for the victim. So, we would obviously *LOVE* to know which countries/actors have access to what other routers (*especially* if it's our routers). Then the question comes down to:

"How would you identify the fact that someone has CNE access to a router?"

(TS//SI//REL) There are a handful of ways that you can tell if anyone has access to **your** router. Like, if you are free to log into your own router, you could frequently log in, run diagnostic commands, pull the IOS (or, more generically, the operating system file), hash that file, and compare that to a list of known-good hashes. But, how would you find out if someone has CNE access to a router that you don't own? How would you identify it if the Chinese had access to a router in Zimbabwe? If all you have is passive network traffic that we've collected in SIGINT, how would **you** do it?

The rest of this post relates to NSA's methods to detect when countries hack routers. We have redacted it to prevent helping those countries improve their ability to hack foreign routers and spy on people undetected.