Full-mesh IPsec network

10 Dos and 500 Don'ts

$ whoami

- Fran Garcia
- SRE @hostedgraphite
- "Break fast and move things"
- Absolutely no networking/cryptography background
- No, seriously, totally unqualified to give this talk

What this talk is not

A success story

An introduction to IPsec

A HOWTO

A set of best practices

What we'll talk about

Hosted Graphite pre-IPsec

What's this IPsec thing anyway and why should I care?

Hosted Graphite does IPsec!

Everything we did wrong (well, the least embarrassing bits)

TL;DW

"IPsec is awful"

(trust me on this one)

Hosted Graphite pre-IPsec

In the beginning, there was n2n...

Early days at Hosted Graphite:

- A way to secure communications for riak was needed
- Not many servers to maintain

Enter n2n:

- P2P VPN software
- Supports compression and encryption
- Really easy to setup and maintain

Wait, so what's the catch?

Best description from an HG engineer: "academic abandonware"

Relies on a central node (supernode):

- No supernode, no network

Single-threaded, not really efficient:

- Became a bottleneck, increasing latency for some of our services

Initially configured on a /24 private IP space

- We were running out of IP addresses!

Replacing n2n

Our requirements:

- Can't depend on fancy networking gear
- Cluster spans multiple locations/providers
- We don't trust the (internal) network!
- Must be efficient enough not to become a bottleneck!
- Simple security model (no complex/dynamic firewall rules)
- Can be implemented reasonably quickly

Potential n2n alternatives

We looked at a bunch of possible alternatives and most of them:

- Were not really designed for a full-mesh network (OpenVPN)
- Encrypt data in user space, incurring a performance penalty (tinc)
- Would tie us to a single provider (like AWS VPCs)
- Involve modifying and rearchitecting all our services
    - (rolling our own application layer encryption)

So after analyzing all our options IPsec won… almost by default

IPsec for the super-impatient

So what's this IPsec thing anyway?

Not a protocol, but a protocol suite

Open standard, which means lots of options for everything

66 RFCs linked from wikipedia page!

What IPsec offers

At the IP layer, it can:

- Encrypt your data (Confidentiality)
- Verify source of received messages (Data-origin authentication)
- Verify integrity of received messages (Data Integrity)

Offers your choice of everything to achieve this

Choices, choices everywhere

What protocol?

- Authentication Header (AH): Just data integrity/authentication*
- Encapsulating Security Payload (ESP): Encryption + integrity/auth (optional)
- AH/ESP

(TL;DR - You probably want ESP)

*Legend says AH only exists to annoy Microsoft

Second choice... Tunnel or Transport mode?

	Encapsulates header	Encapsulates payload	Works for host-to-host	Works for site-to-site
Tunnel Mode	YES	YES	YES	YES
Transport Mode	NO	YES	YES	NO

*Transport mode might incur in a slightly smaller overhead and be a bit simpler to set up

IPsec: What's a SP (security policy)?

Consulted by the kernel when processing traffic (inbound and outbound)

```
"From host A to host B use ESP in transport mode"

"From host C to host D's port 443 do not use IPsec at all"
```

Stored in the SPD (Security Policy Database) inside the kernel

IPsec: What's a SA (Security Association)?

Secured **unidirectional** connection between peers:

- So **need two** for bidirectional communication (hosta->hostb, hostb->hosta)

Contains keys and other attributes like its lifetime, IP address of peer...

Stored in the SAD (Security Association Database) inside the kernel

IKE? Who's IKE?

"Internet Key Exchange"

Negotiate algorithms/keys needed to establish secure channel between peers

A key management daemon does it in user space, consists of 2 phases

IPsec: IKE Phase 1

Lives inside the key management daemon (in user space)

Hosts negotiate proposals on how to authenticate and secure the channel

Negotiated session keys used to establish actual (multiple) IPsec SAs later
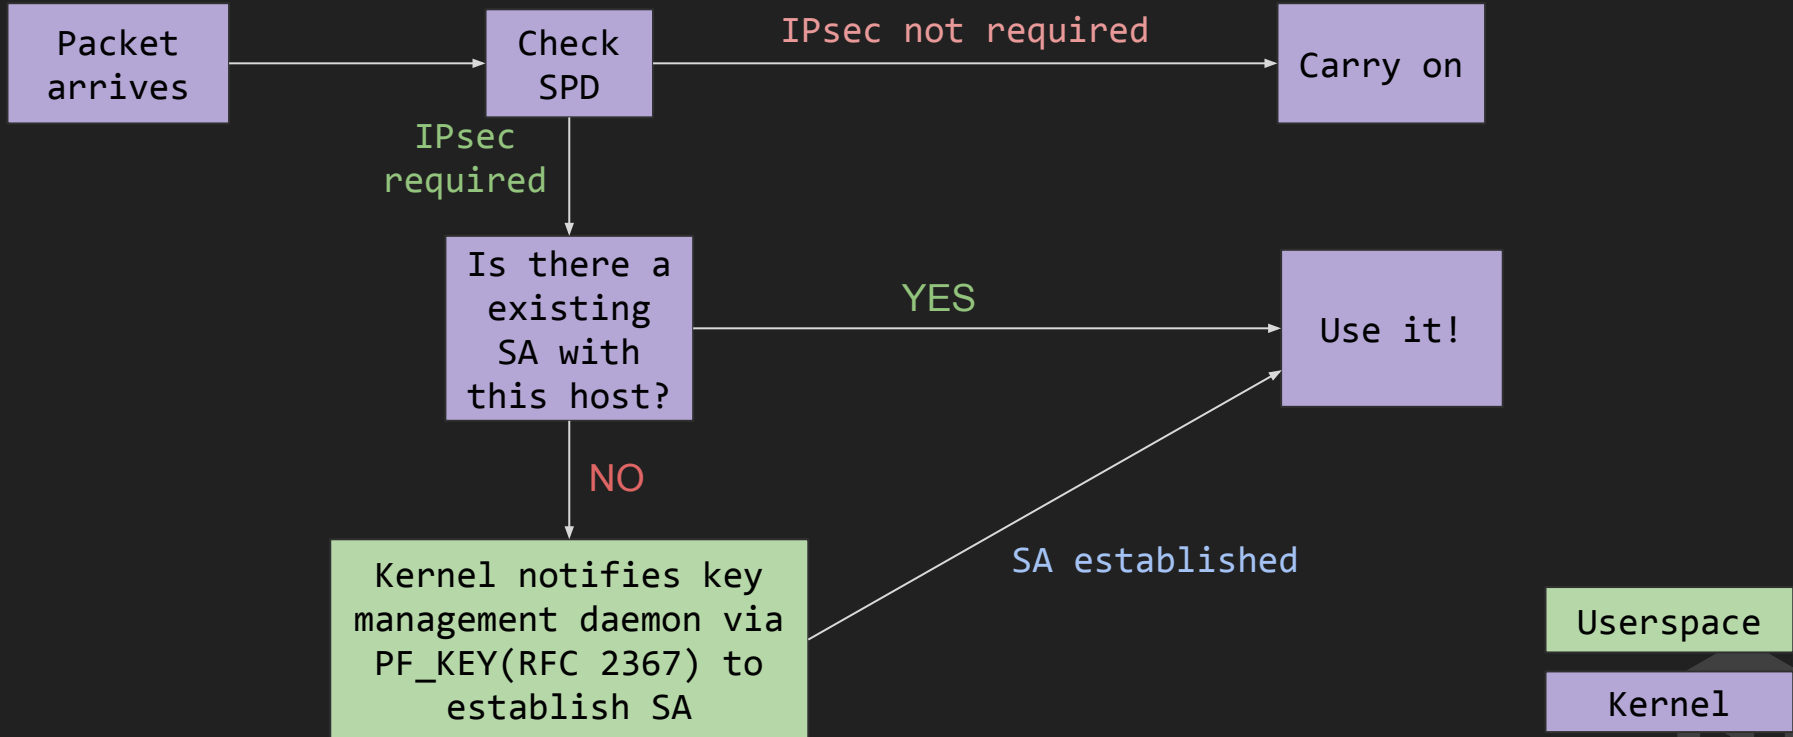
IPsec: Phase 2

Negotiates IPsec SA parameters (protected by IKE SA) using phase 1 keys

Establishes the actual IPsec SA (and stores it in SADB)

Can renegotiate when close to end of lifetime

Life of an IPsec packet

Packet arrives → Check SPD

Check SPD → **IPsec not required** → Carry on

Check SPD → **IPsec required** → Is there a existing SA with this host?

Is there a existing SA with this host? → **YES** → Use it!

Is there a existing SA with this host? → **NO** → Kernel notifies key management daemon via PF_KEY(RFC 2367) to establish SA

Kernel notifies key management daemon via PF_KEY(RFC 2367) to establish SA → **SA established** → Use it!

Userspace

Kernel

Some helpful commands

`ip xfrm` is pretty powerful. Some basics:

```
$ ip xfrm policy          # Dump the contents of the SPDB

$ ip xfrm state           # Dump the contents of the SADB

$ ip xfrm monitor         # Dump all changes to SADB and SPDB as they happen

$ ip xfrm state flush     # Flush all state in the SADB (dangerous!)
```

Documentation is... not great: http://man7.org/linux/man-pages/man8/ip-xfrm.8.html

So what has IPsec ever done for us?

Encryption happens inside the kernel, so it's fast!

Using the right algorithms/settings it can be fairly secure

It's a standard, so there are good practices to use it securely

Very flexible, which is useful if you have:

- Hardware distributed across different datacenters/providers
- No real control over your network infrastructure

Hosted Graphite does IPsec!

Our migration: n2n -> IPsec

Big time constraints: n2n was unreliable and preventing us from scaling

We had trouble finding reports of people using IPsec in the same way*…

...So we had to improvise a bit.

After careful planning and testing we rolled it out to our production cluster...

* Notable exception: pagerduty's Doug Barth at Velocity 2015 http://conferences.oreilly.
com/velocity/devops-web-performance-2015/public/schedule/detail/41454

WORST

MIGRATION

EVER

WORST. MIGRATION. EVER

Migration attempt resulted in multi-day incident:

http://status.hostedgraphite.com/incidents/gw2v1rhm8p5g

Took two days to stabilize, a full week to resolve the incident.

Lots of issues not found during testing

n2n -> IPsec migration aftermath

Back to drawing board, came up with another plan

Spent almost 3 months slowly rolling it out and fixing bugs:

- Also known as "the worst three months of my life"
- Big team effort, everybody pitched in

Still worth it, things are stable now and we've learned a lot

Our stack: present day

Our IPsec stack: present day

Hundreds of hosts using ESP in transport mode (full-mesh)

Several clusters, isolated from each other

Using ipsec-tools with racoon as key management daemon

Our config: iptables

```
# Accept all IKE traffic, also allowing NAT Traversal (UDP 4500)
-A ufw-user-input -p udp --dport 500 -j ACCEPT
-A ufw-user-input -p udp --dport 4500 -j ACCEPT



# Allow  all ESP traffic, if it has a formed IPsec SA we trust it
-A ufw-user-input -p esp -j ACCEPT
```

Our config: Security Policies (`/etc/ipsec-tools.conf`)

Node1 = 1.2.3.4   Node2 = 5.6.7.8

On node1:

```
# require use of IPsec for all other traffic with node2
spdadd 1.2.3.4 5.6.7.8 any -P out ipsec esp/transport//require;
spdadd 5.6.7.8 1.2.3.4 any -P in ipsec esp/transport//require;
```


On node2:

```
# require use of IPsec for all other traffic with node1
spdadd 5.6.7.8 1.2.3.4 any -P out ipsec esp/transport//require;
spdadd 1.2.3.4 5.6.7.8 any -P in ipsec esp/transport//require;
```

Our config: Security Policies (`/etc/ipsec-tools.conf`)

What about management hosts?

Node1 = 1.2.3.4   PuppetMaster = 5.6.7.8

On node1:

```
# Only require IPsec for port 8140 on the puppet master
spdadd 1.2.3.4 5.6.7.8[8140] any -P out ipsec esp/transport//require;
spdadd 5.6.7.8[8140] 1.2.3.4 any -P in ipsec esp/transport//require;
```

Everything else will get dropped by the firewall

Our config: Security Policies (`/etc/ipsec-tools.conf`)

```
# Exclude ssh traffic:
spdadd 0.0.0.0/0[22] 0.0.0.0/0 tcp -P in prio def +100 none;
spdadd 0.0.0.0/0[22] 0.0.0.0/0 tcp -P out prio def +100 none;
spdadd 0.0.0.0/0 0.0.0.0/0[22] tcp -P in prio def +100 none;
spdadd 0.0.0.0/0 0.0.0.0/0[22] tcp -P out prio def +100 none;


# Exclude ICMP traffic (decouple ping and the like from IPsec):
spdadd 0.0.0.0/0 0.0.0.0/0 icmp -P out prio def +100 none;
spdadd 0.0.0.0/0 0.0.0.0/0 icmp -P in prio def +100 none;
```

Our config: racoon (`/etc/racoon.conf`)

Phase 1:

```
remote anonymous {
    exchange_mode main;
    dpd_delay 0;
    lifetime time 24 hours;
    nat_traversal on;
    proposal {
        authentication_method pre_shared_key;
        dh_group modp3072;
        encryption_algorithm aes;
        hash_algorithm sha256;
    }
}
```

Our config: racoon (`/etc/racoon.conf`)

Phase 2:

```
sainfo anonymous {
    pfs_group modp3072;
    encryption_algorithm aes;
    authentication_algorithm hmac_sha256;
    compression_algorithm deflate;

    lifetime time 8 hours;
}
```

10 DOS AND 500 DONT'S

Disclaimer:

(We don't really have 10 dos)

Don't use ipsec-tools/racoon! (like we did)

Not actively maintained (Last release on early 2014)

Buggy

But the only thing that worked for us under time/resource constraints

LibreSwan seems like a strong alternative

"The mystery of the disappearing SAs"

Some hosts unable to establish SAs on certain cases

Racoon would complain of SAs not existing (kernel would disagree):

```
ERROR: no policy found: id:281009.
```

racoon's internal view of the SADB would get out of sync with the kernel's

We suspect corruption in racoon's internal state for the SADB

"The mystery of the disappearing SAs"

Restarting racoon fixes it, but that wipes out all your SAs!

Workaround: Force racoon to reload both SADB and config

```
killall -HUP racoon
```


Forcing periodic reloads prevents the issue from reoccurring ¯_(ツ)_/¯

Don't blindly force all traffic to go through IPsec

Account for everything that needs an exception:

- SSH, ICMP, etc

You'll need to be able to answer these two questions:

- "Is the network broken?"
- "Is IPsec broken?"

"Yo dawg, I heard you like encrypted traffic…"

If migrating from an existing VPN, make sure to exclude it from IPsec traffic

During our initial rollout our SPs forced our n2n traffic through IPsec…

… Which still wasn't working reliably enough…

… Effectively killing our whole internal network

Don't just enable DPD… without testing

What's DPD?

- DPD: Dead Peer Detection (RFC3706)
- Liveness checks on Phase 1 relationships
- If no response to R-U-THERE clears phase 1 and 2 relationships…

Sounds useful but test it in your environment first:

- racoon implementation is buggy!

"The trouble with DPDs"

In our case, enabling DPD results in 100s of SAs between two hosts:

- Every failed DPD check resulting in extra SAs

Combination of factors:

- Unreliable network
- Bugs in racoon

We ended up giving up on DPD

Don't just disable DPD either

DPD can be legitimately useful

Example: What happens when rebooting a host?

Other nodes might not realise their SAs are no longer valid!

DPD: Rebooting hosts

bender's SAD:

5.6.7.8 -> 1.2.3.4 (spi: 0x01)

1.2.3.4 -> 5.6.7.8 (spi: 0x02)

flexo's SAD:

5.6.7.8 -> 1.2.3.4 (spi: 0x01)

1.2.3.4 -> 5.6.7.8 (spi: 0x02)

These are two happy hosts right now…

bender -> flexo (using spi 0x02) traffic is received by flexo

flexo -> bender (using spi 0x01) traffic is received by bender!

… But let's say we reboot bender!

DPD: Rebooting hosts

bender's SAD:

~~5.6.7.8 -> 1.2.3.4 (spi: 0x02)~~

~~1.2.3.4 -> 5.6.7.8 (spi: 0x01)~~

flexo's SAD:

5.6.7.8 -> 1.2.3.4 (spi: 0x02)

1.2.3.4 -> 5.6.7.8 (spi: 0x01)

bender's SADB is now empty

flexo->bender traffic (using spi 0x02) will be broken until:

- bender->flexo traffic forces establishment of new SAs
- The SAs on flexo's side expire

DPD: Just roll your own

Our solution: Implement our own phase 2 liveness check

Check a known port for every host we have a mature SA with:

- Clear the SAs if ${max_tries} timeouts

Bonus points: Also check a port that won't use IPsec to compare

Do instrument all the things!

You'll ask yourself "is the network broken or just IPsec?" a lot
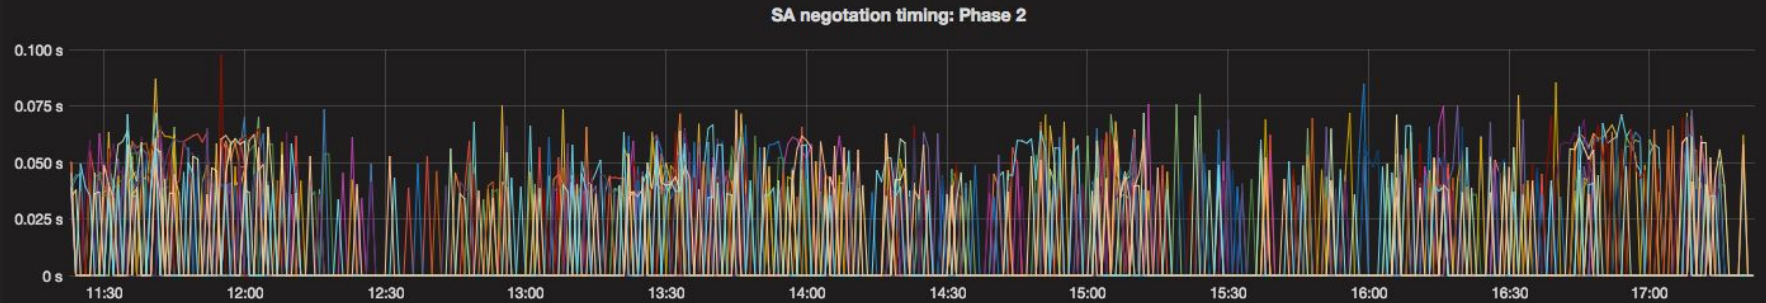
So better have lots of data!

Built racoon to emit timing info on logs (build with `--enable-stats`)

Diamond collector gathers and send metrics from:

- racoon logs
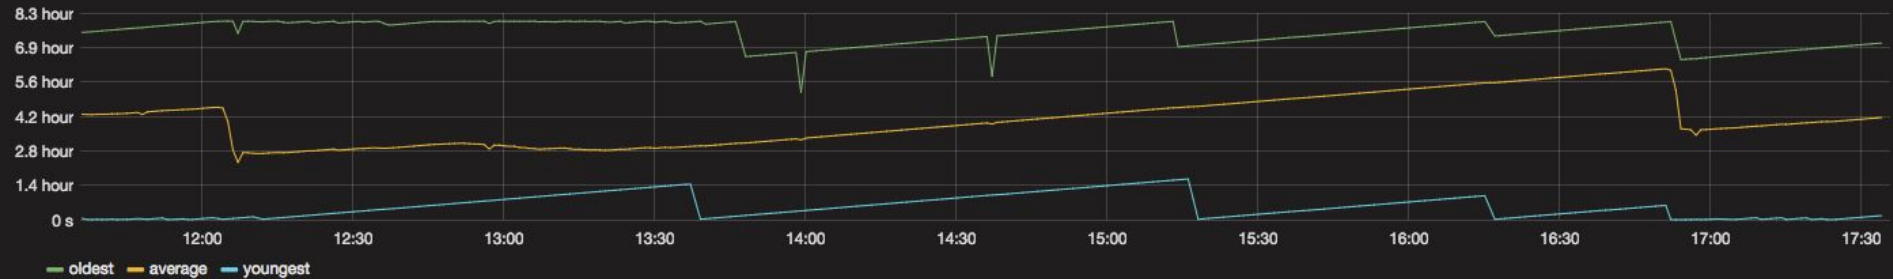- SADB

Instrumenting the racoon logs

Instrumenting the SADB

SA state

Legend: larval, mature, dying

SA age

Legend: oldest, average, youngest

Do instrument all the things!

Kernel metrics also useful (if available!)

You want your kernel compiled with `CONFIG_XFRM_STATISTICS`

```
$ cat /proc/net/xfrm_stat
XfrmInError                    0
XfrmInBufferError              0
...
```

(Very) brief descriptions: https://www.kernel.org/doc/Documentation/networking/xfrm_proc.txt

Instrumenting kernel xfrm stats: XfrmInNoStates

XfrmInNoStates: Times we've received data for an SA we know nothing about

"The case of the sad halfling"

bender's SAD:

```
5.6.7.8 -> 1.2.3.4 (spi: 0x02)

1.2.3.4 -> 5.6.7.8 (spi: 0x01)
```

flexo's SAD:

```
5.6.7.8 -> 1.2.3.4 (spi: 0x02)

1.2.3.4 -> 5.6.7.8 (spi: 0x01)
```

These are two happy hosts right now...

… But let's say one SA "disappears" during a brief netsplit:

```
bender$ echo "deleteall 5.6.7.8 1.2.3.4 esp ; " | setkey -c

# The 5.6.7.8 1.2.3.4 association gets removed from bender
```

"The case of the sad halfling"

bender's SAD:

~~5.6.7.8 -> 1.2.3.4 (spi: 0x02)~~

1.2.3.4 -> 5.6.7.8 (spi: 0x01)

flexo's SAD:

5.6.7.8 -> 1.2.3.4 (spi: 0x02)

1.2.3.4 -> 5.6.7.8 (spi: 0x01)

Any communication attempt will fail!

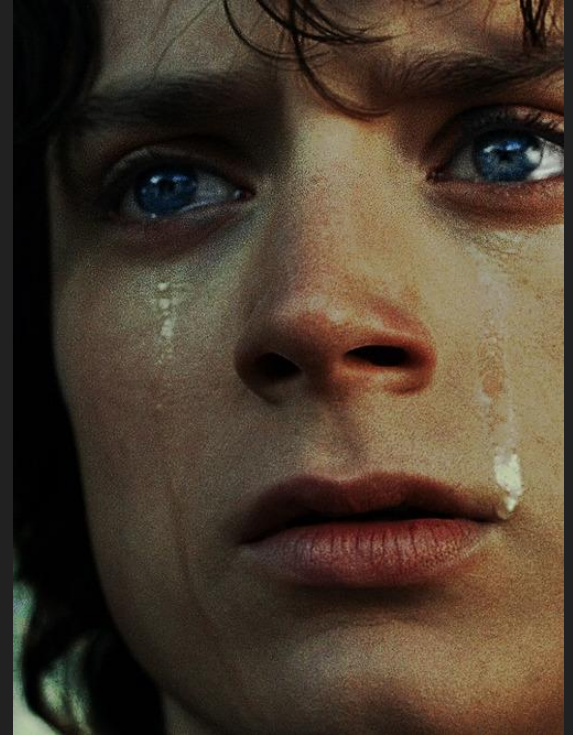bender -> flexo (using spi 0x01) traffic is received by flexo

flexo -> bender (using spi 0x02) traffic is ignored by bender!

"The case of the sad halfling"

We built a custom daemon for detecting it

Highlights need for:

- Phase 2 liveness checks
- Metrics for everything!

Don't flush/restart on changes

Never restart racoon!

A racoon restart will flush all phase 1 and 2 SAs:

- Negotiating ~1000 SAs at once is no fun

To flush an individual SA: `ip xfrm state delete`

To reload config changes: `killall -HUP racoon`

Don't flush/restart on changes

When adding/removing hosts, do not flush both SPD and SAD!

Just flush/reload your SPD and let unwanted SAs expire

SAs not reflected in the SPD will never get used

Can flush that SA individually if feeling paranoid

Can just include `spdflush;` in your `ipsec-tools.conf` and reload with:

`setkey -f /etc/ipsec-tools.conf`

You don't have the same tools available

tcpdump will just show ESP traffic, not its content:

```
15:47:51.511135 IP 1.2.3.4 > 5.6.7.8: ESP(spi=0x00fb0c52,seq=0x1afa), length 84

15:47:51.511295 IP 5.6.7.8 > 1.2.3.4: ESP(spi=0x095e5523,seq=0x173a), length 84
```

Traffic can be decrypted with wireshark/tshark if you dump the keys first

You don't have the same tools available

Can use tcpdump with netfilter logging framework:

```
$ iptables -t mangle -I PREROUTING -m policy --pol ipsec --dir in -j NFLOG --nflog-group 5

$ iptables -t mangle -I POSTROUTING -m policy --pol ipsec --dir out -j NFLOG --nflog-group 5

$ tcpdump -i nflog:5
```

Doesn't allow most filters

Might need to increase the buffer size

You don't have the same tools available

Traceroute will attempt to use udp by default:

```
$ traceroute that.other.host
traceroute to that.other.host (5.6.7.8), 30 hops max, 60 byte packets
 1  * * *
 2  * * *
 3  * * *
 4  that.other.host (5.6.7.8)  0.351 ms  0.295 ms  0.297 ms
```

You can force it to use ICMP with `traceroute -I`

Do use certs for auth, or don't use a weak PSK

PSK is great for getting started if you don't have PKI in place (we didn't)

But please:

- Use a strong PSK (if you must use PSK)
- Enable PFS (Perfect Forward Secrecy)
- Do not use aggressive mode for phase 1

Not following all that makes the NSA happy!

Don't trust the (kernel) defaults!

Careful with `net.ipv4.xfrm4_gc_thresh`

Associations might be garbage collected before they can succeed!

If `3.6 > $(uname -r) < 3.13`:

Default `(1024)` might be too low

GC will cause performance issues

Can refuse new allocations if you hit `(1024 * 2)` dst entries

Don't trust the defaults!

Beware of IPsec overhead and MTU/MSS:

A 1450 bytes IP packet becomes:

- 1516 bytes in Transport mode
- 1532 bytes in Tunnel mode

(More if enabling NAT Traversal!)

Path MTU Discovery should help, but test it first!

Thanks!

Questions?

Hated it? Want to say hi?

fran@hostedgraphite.com

@hostedgraphite

hostedgraphite.com/jobs