KEEPING YOUR RACK COOL WITH ONE "/IP ROUTE RULE"

Marek Isalski – marek @ faelix.net – @maznu
faelix limited – https://faelix.net/ – @faelix

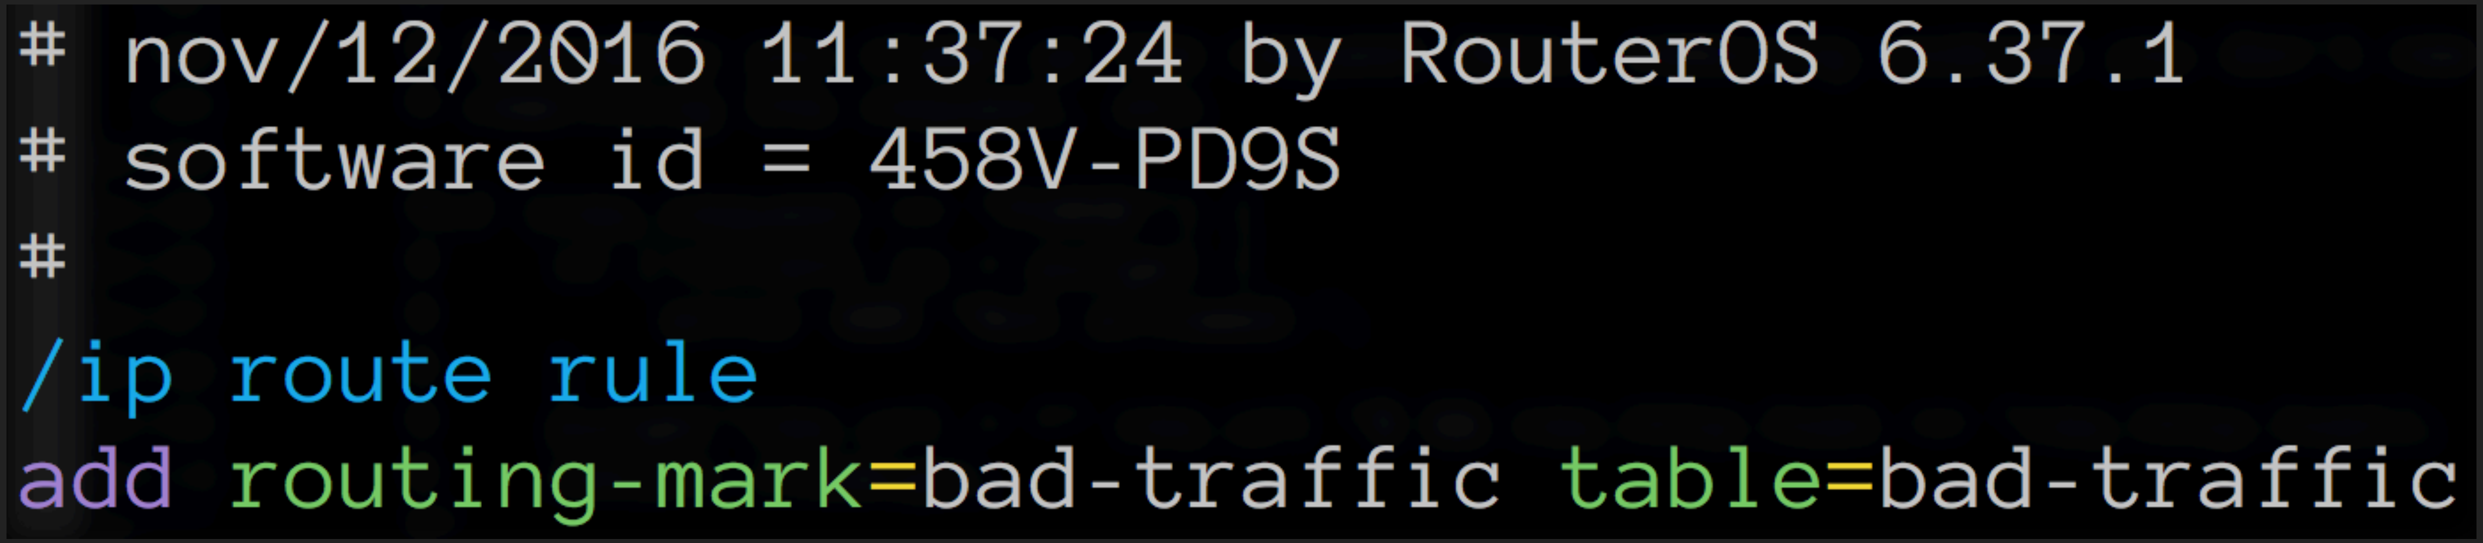

```
# nov/12/2016 11:37:24 by RouterOS 6.37.1
# software id = 458V-PD9S
#
/ip route rule
add routing-mark=bad-traffic table=bad-traffic
```

THE END

QUESTIONS ETC?

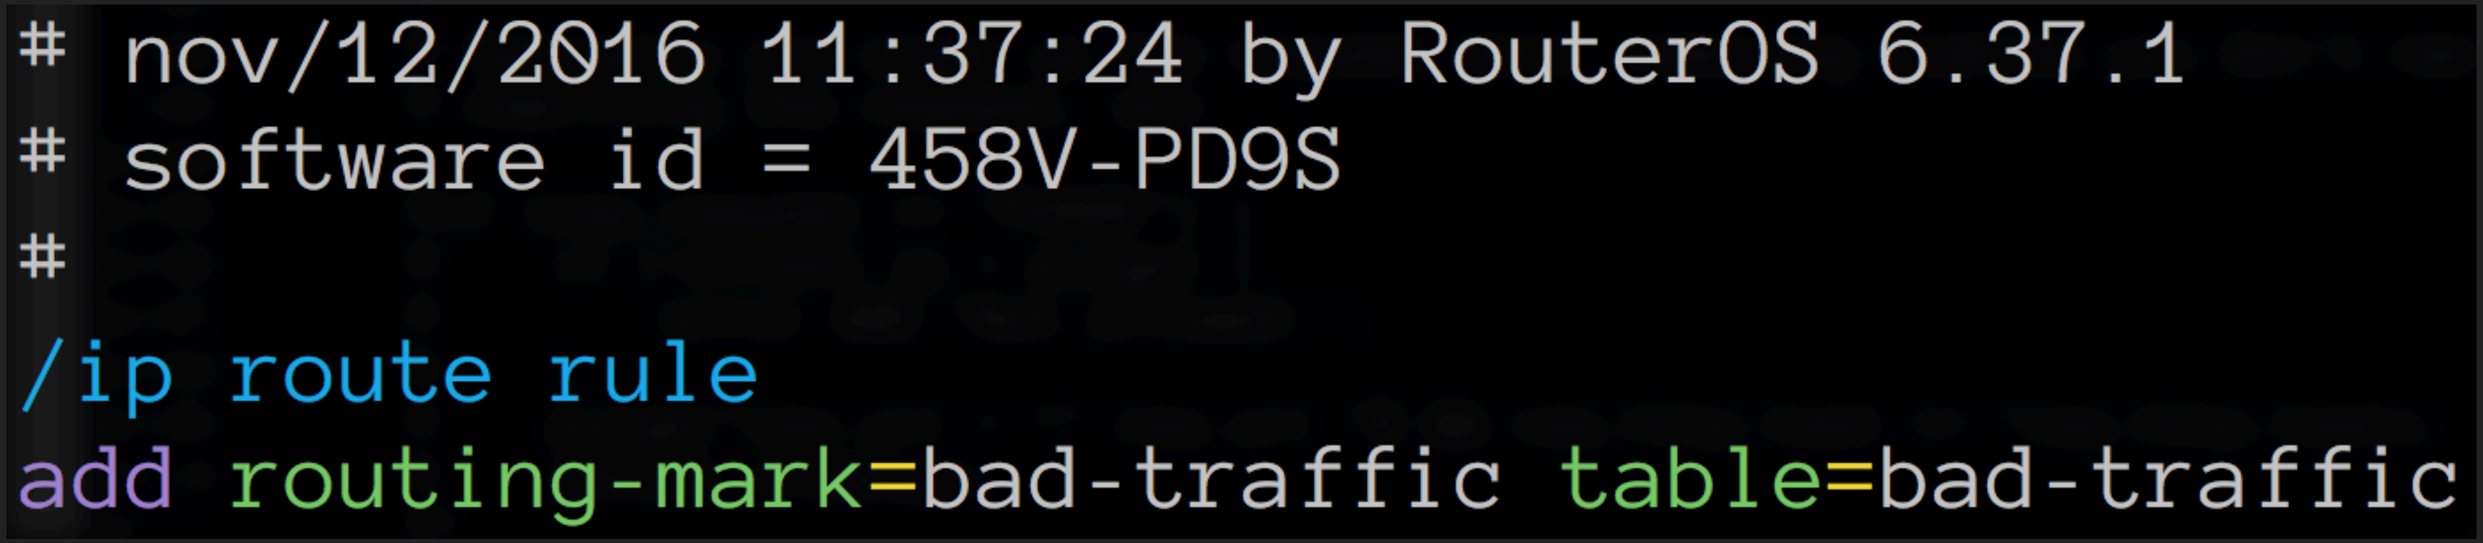

```
# nov/12/2016 11:37:24 by RouterOS 6.37.1
# software id = 458V-PD9S
#
/ip route rule
add routing-mark=bad-traffic table=bad-traffic
```

NOT SO FAST…

;-)

HOW FAELIX ARRIVED AT THIS IDEA

- Part 1:
  - About our network and what we do
  - Our experience using MikroTik at the provider edge
- Part 2:
  - Zero filter rules! :-)

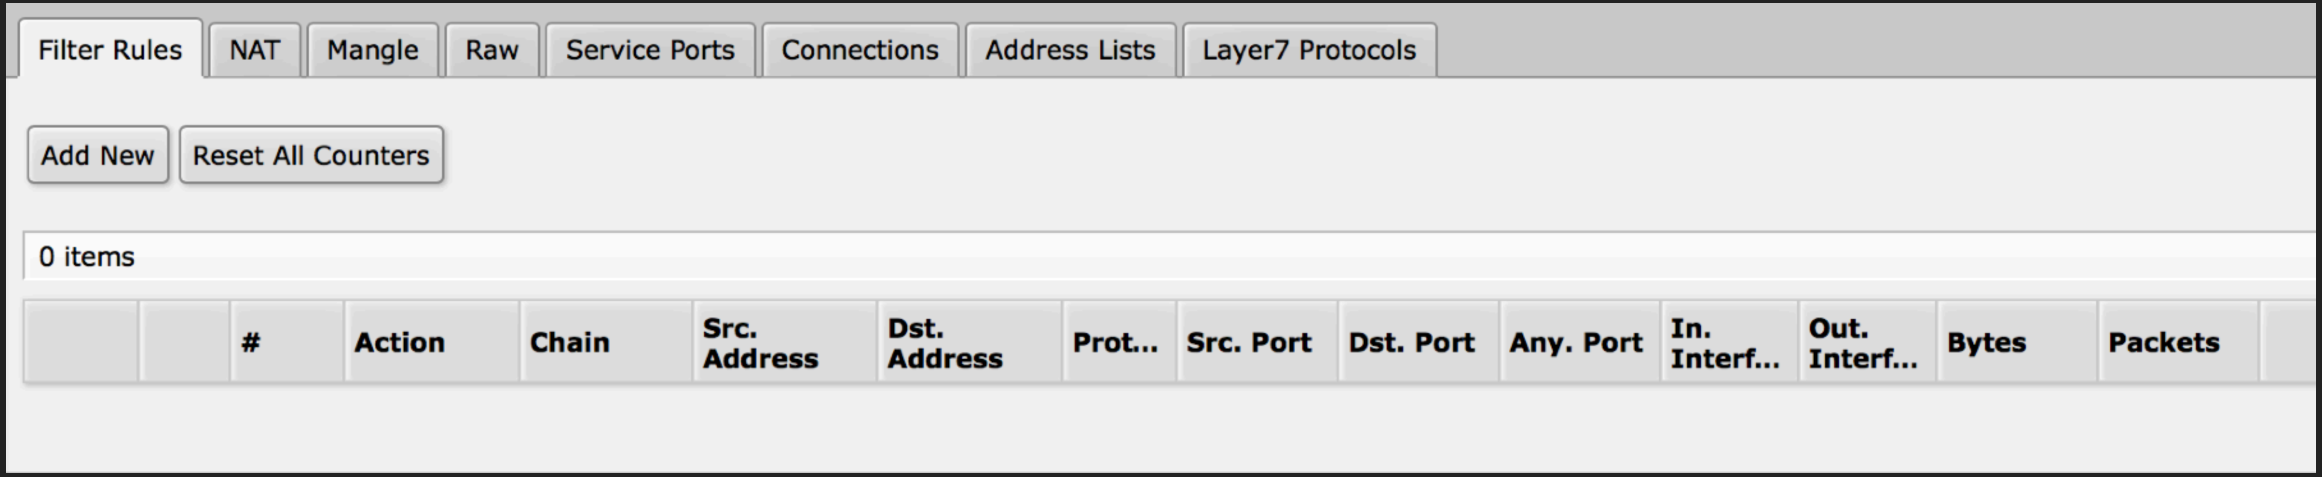

0 items

		#	Action	Chain	Src. Address	Dst. Address	Prot...	Src. Port	Dst. Port	Any. Port	In. Interf...	Out. Interf...	Bytes	Packets

PART 1:

MIKROTIK AT THE PROVIDER EDGE

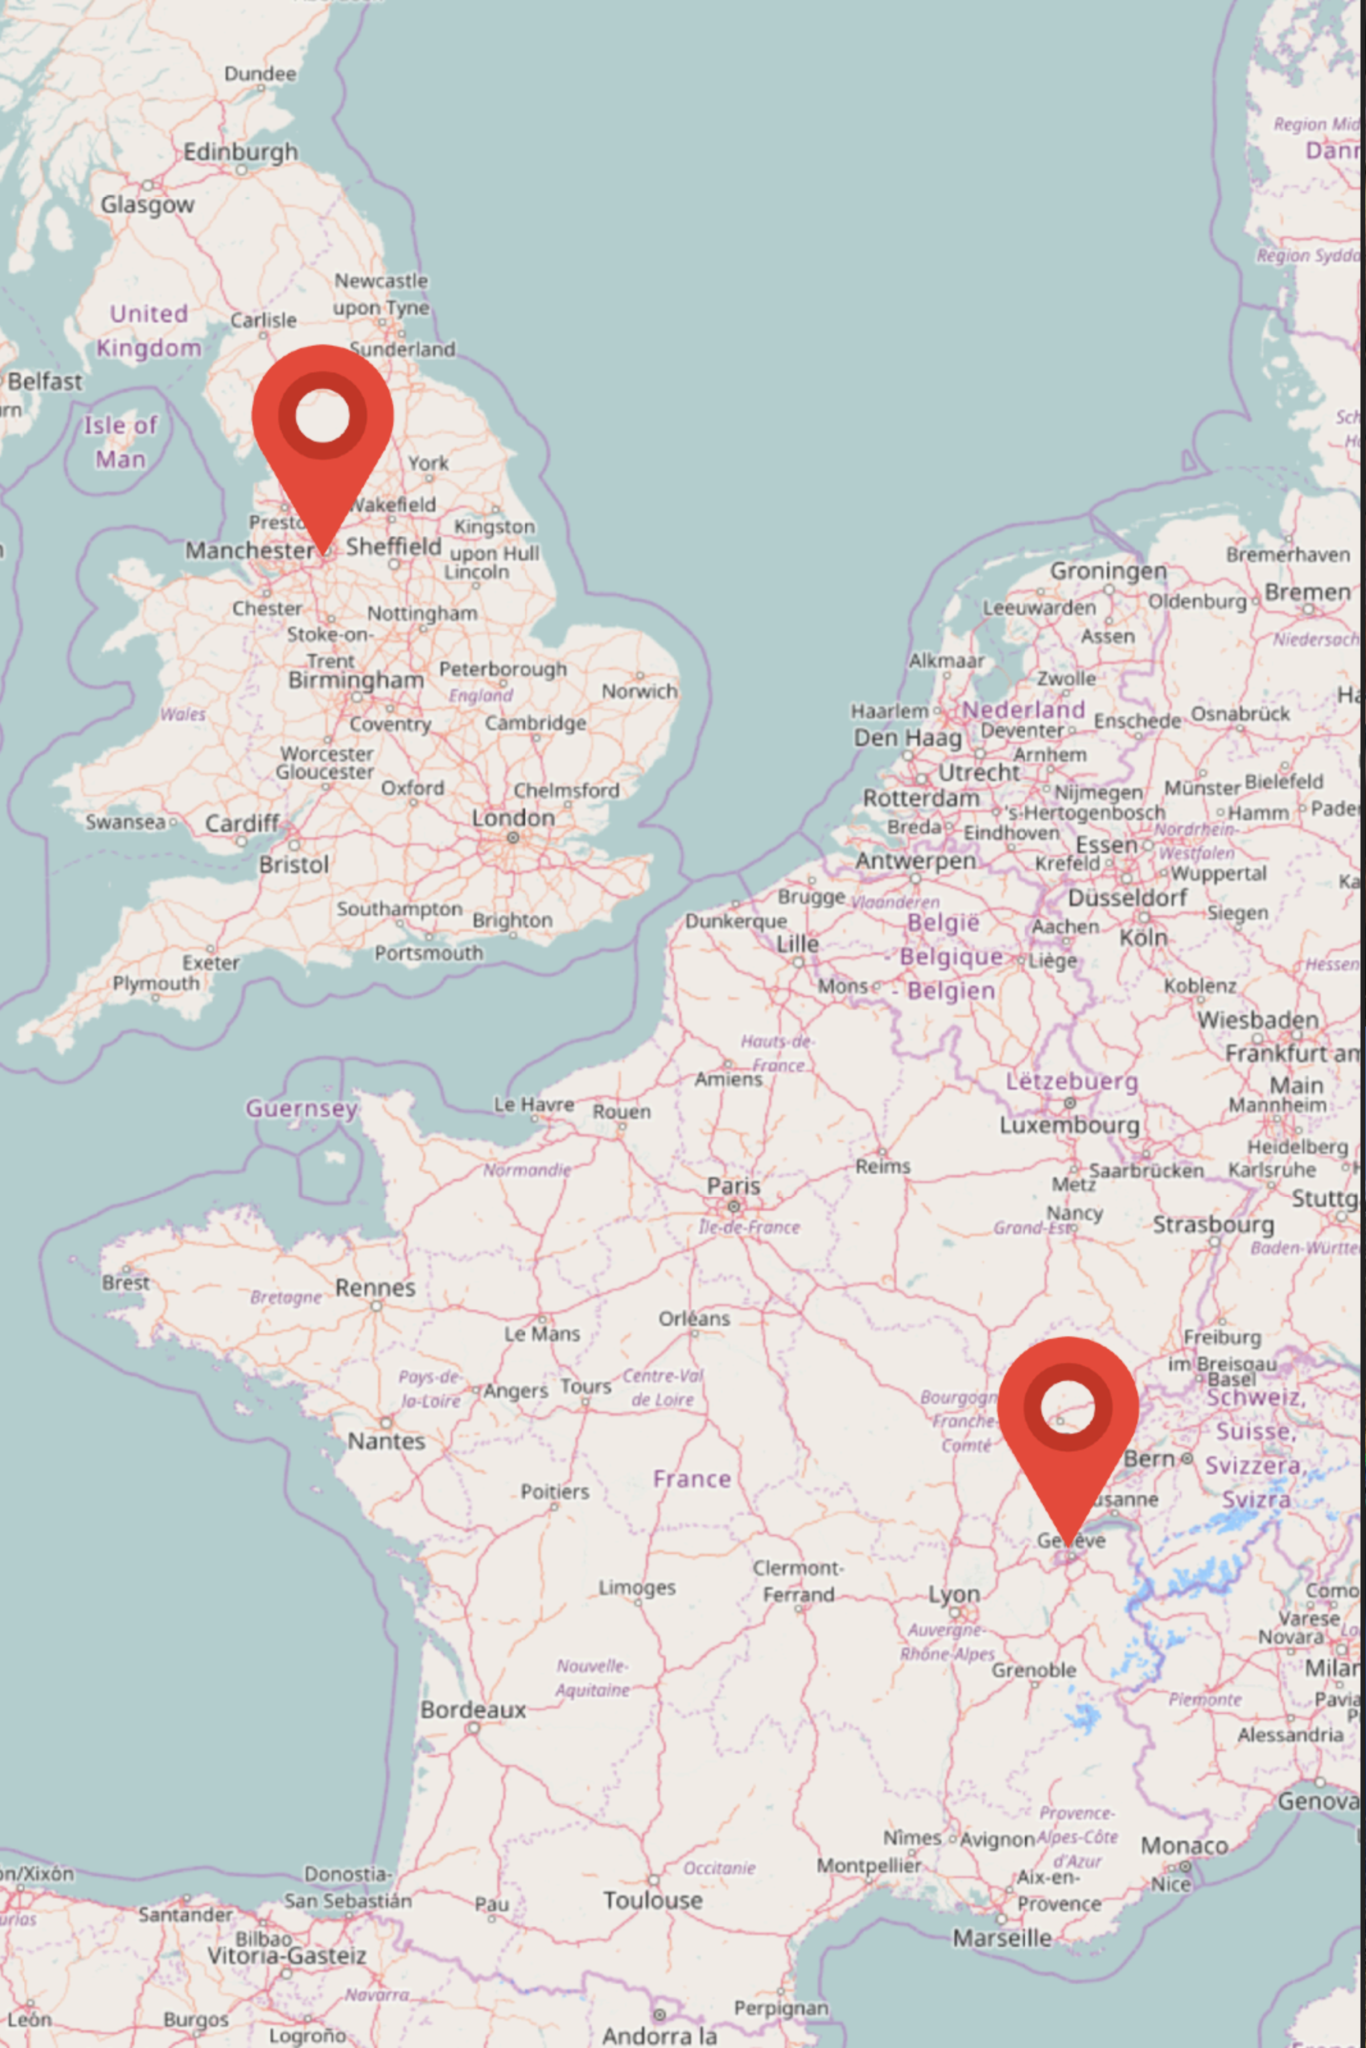
Dundee
Edinburgh
Glasgow
United Kingdom
Carlisle
Newcastle upon Tyne
Sunderland
Belfast
Isle of Man
York
Wakefield
Preston
Manchester
Sheffield
Kingston upon Hull
Lincoln
Chester
Nottingham
Stoke-on-Trent
Birmingham
Peterborough
England
Norwich
Wales
Coventry
Cambridge
Worcester
Gloucester
Oxford
Chelmsford
Swansea
Cardiff
London
Bristol
Southampton
Brighton
Portsmouth
Exeter
Plymouth
Guernsey
Region Midtjylland
Danmark
Region Syddanmark
Groningen
Bremerhaven
Leeuwarden
Oldenburg
Bremen
Assen
Niedersachsen
Alkmaar
Zwolle
Haarlem
Nederland
Enschede
Osnabrück
Den Haag
Deventer
Arnhem
Utrecht
Rotterdam
Nijmegen
Münster
Bielefeld
Breda
's-Hertogenbosch
Hamm
Eindhoven
Nordrhein-Westfalen
Essen
Antwerpen
Krefeld
Wuppertal
Brugge
Vlaanderen
Düsseldorf
Siegen
Dunkerque
België
Aachen
Köln
Lille
Belgique
Liège
Mons
Belgien
Koblenz
Hessen
Hauts-de-France
Wiesbaden
Frankfurt am Main
Amiens
Lëtzebuerg
Mannheim
Luxembourg
Le Havre
Rouen
Heidelberg
Normandie
Reims
Saarbrücken
Karlsruhe
Paris
Metz
Stuttgart
Île-de-France
Nancy
Grand-Est
Strasbourg
Baden-Württemberg
Brest
Rennes
Bretagne
Freiburg im Breisgau
Le Mans
Orléans
Basel
Pays-de-la-Loire
Angers
Tours
Centre-Val de Loire
Bourgogne-Franche-Comté
Schweiz, Suisse, Svizzera, Svizra
Nantes
Bern
France
Lausanne
Poitiers
Genève
Limoges
Clermont-Ferrand
Lyon
Como
Varese
Novara
Auvergne-Rhône-Alpes
Grenoble
Milano
Nouvelle-Aquitaine
Pavia
Bordeaux
Piemonte
Alessandria
Genova
Nîmes
Avignon
Provence-Alpes-Côte d'Azur
Monaco
Montpellier
Nice
Occitanie
Aix-en-Provence
Donostia-San Sebastián
Santander
Pau
Toulouse
Marseille
Bilbao
Vitoria-Gasteiz
Navarra
Perpignan
León
Burgos
Logroño
Andorra la Vella

ABOUT FAELIX

- Mostly-hosting ISP
- Security, social issues, environment
- Based in Manchester, UK = local footprint
- ≈50% of servers in Geneva, CH = excellent energy efficiency
- Multi-homed, multi-site, autonomous system: AS41495

EYEBALLS vs CONTENT

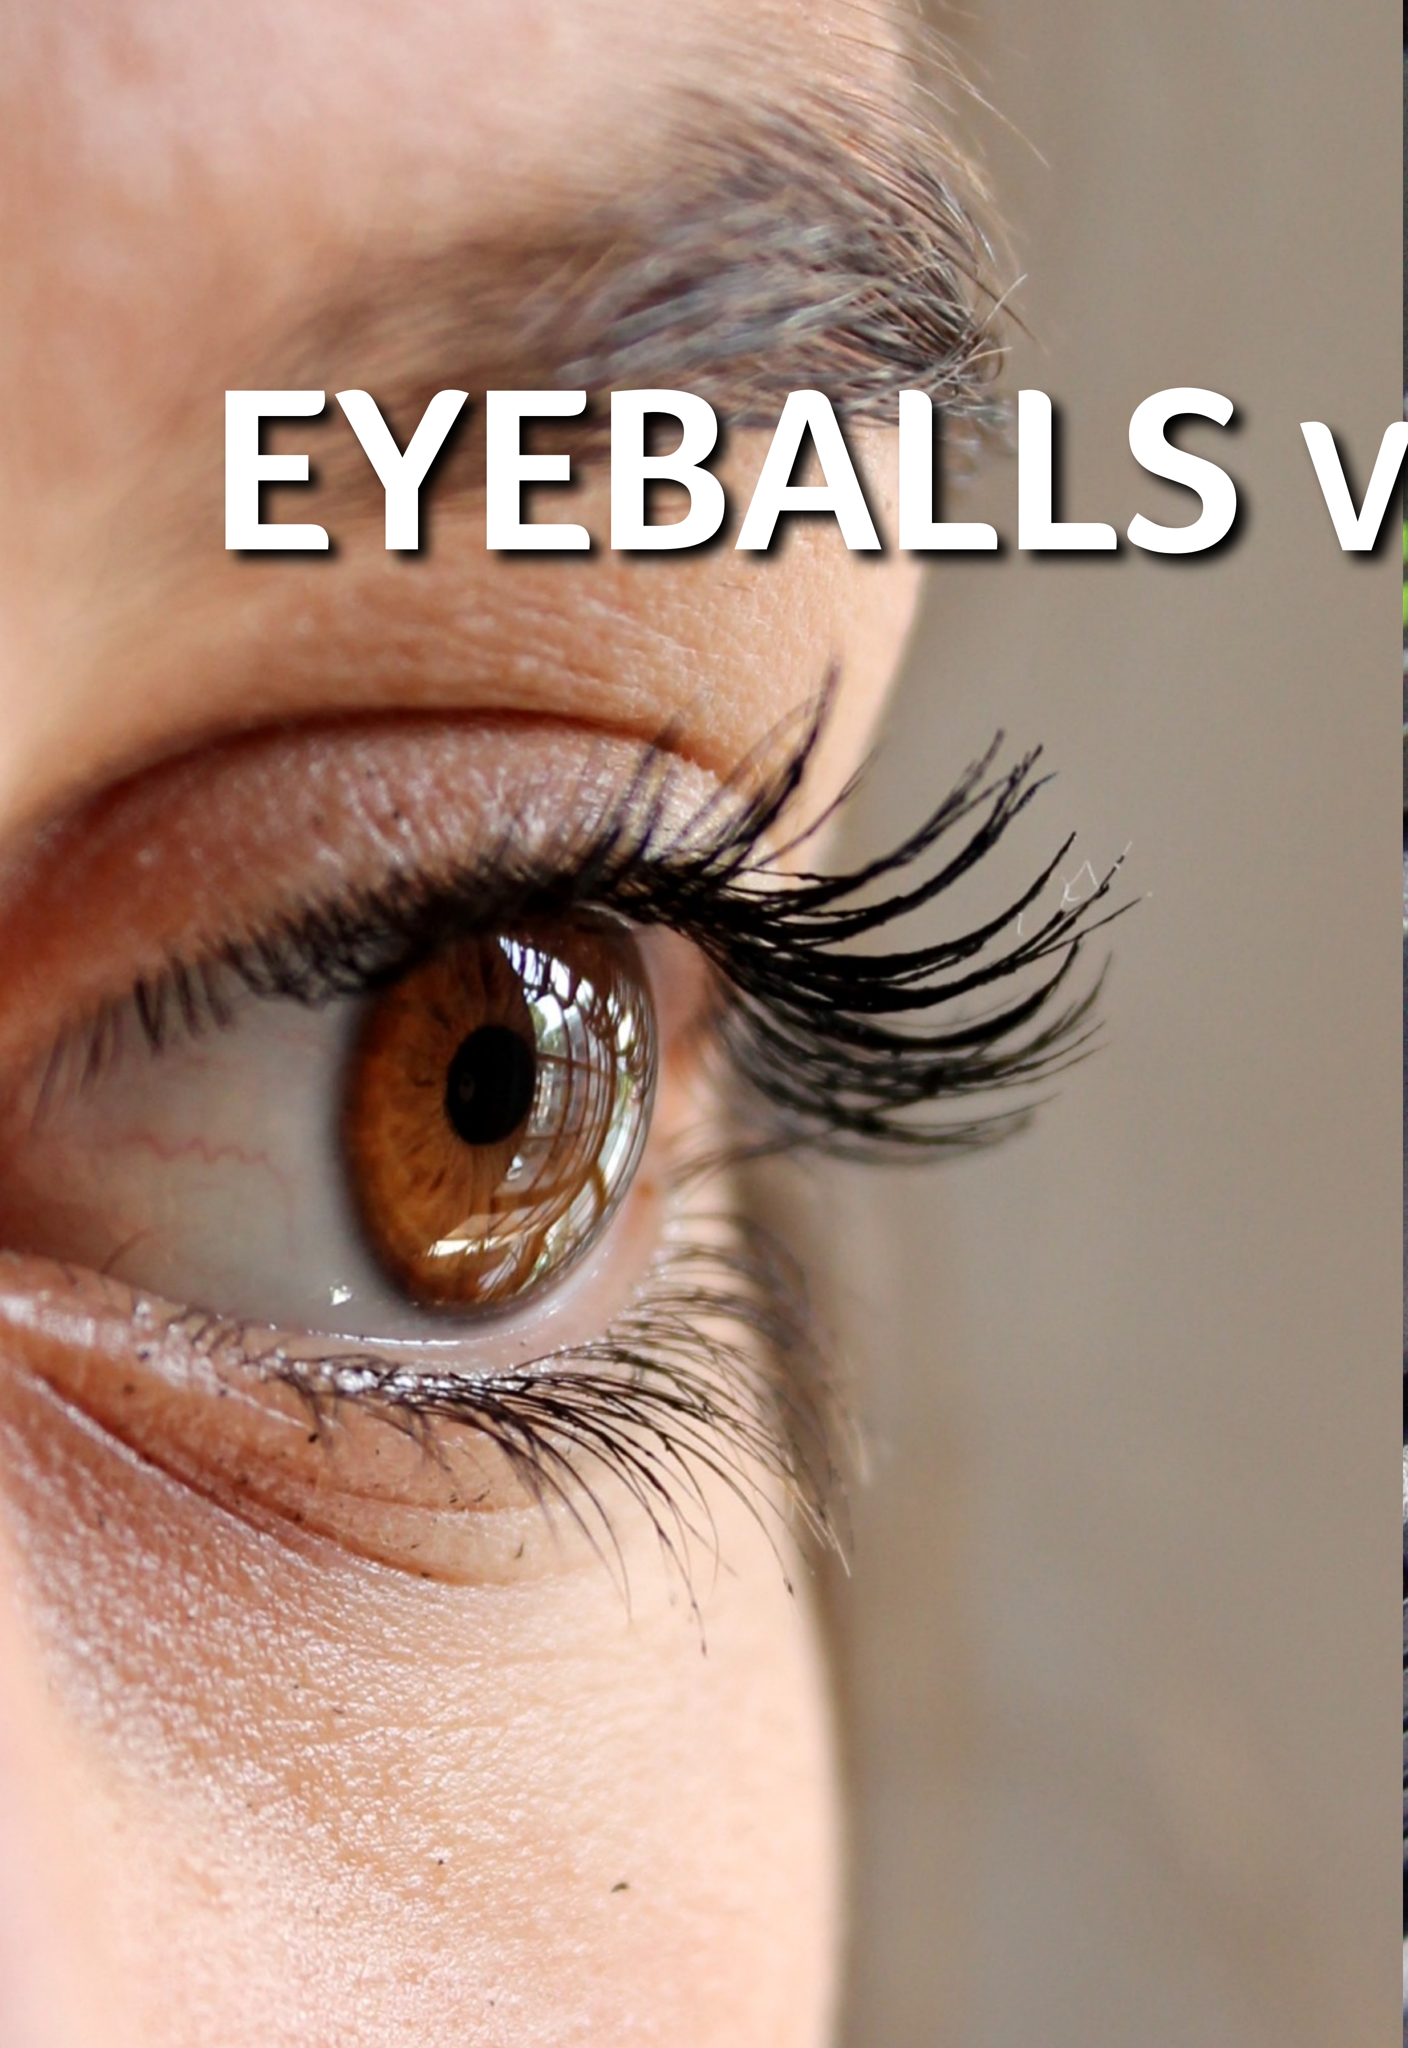

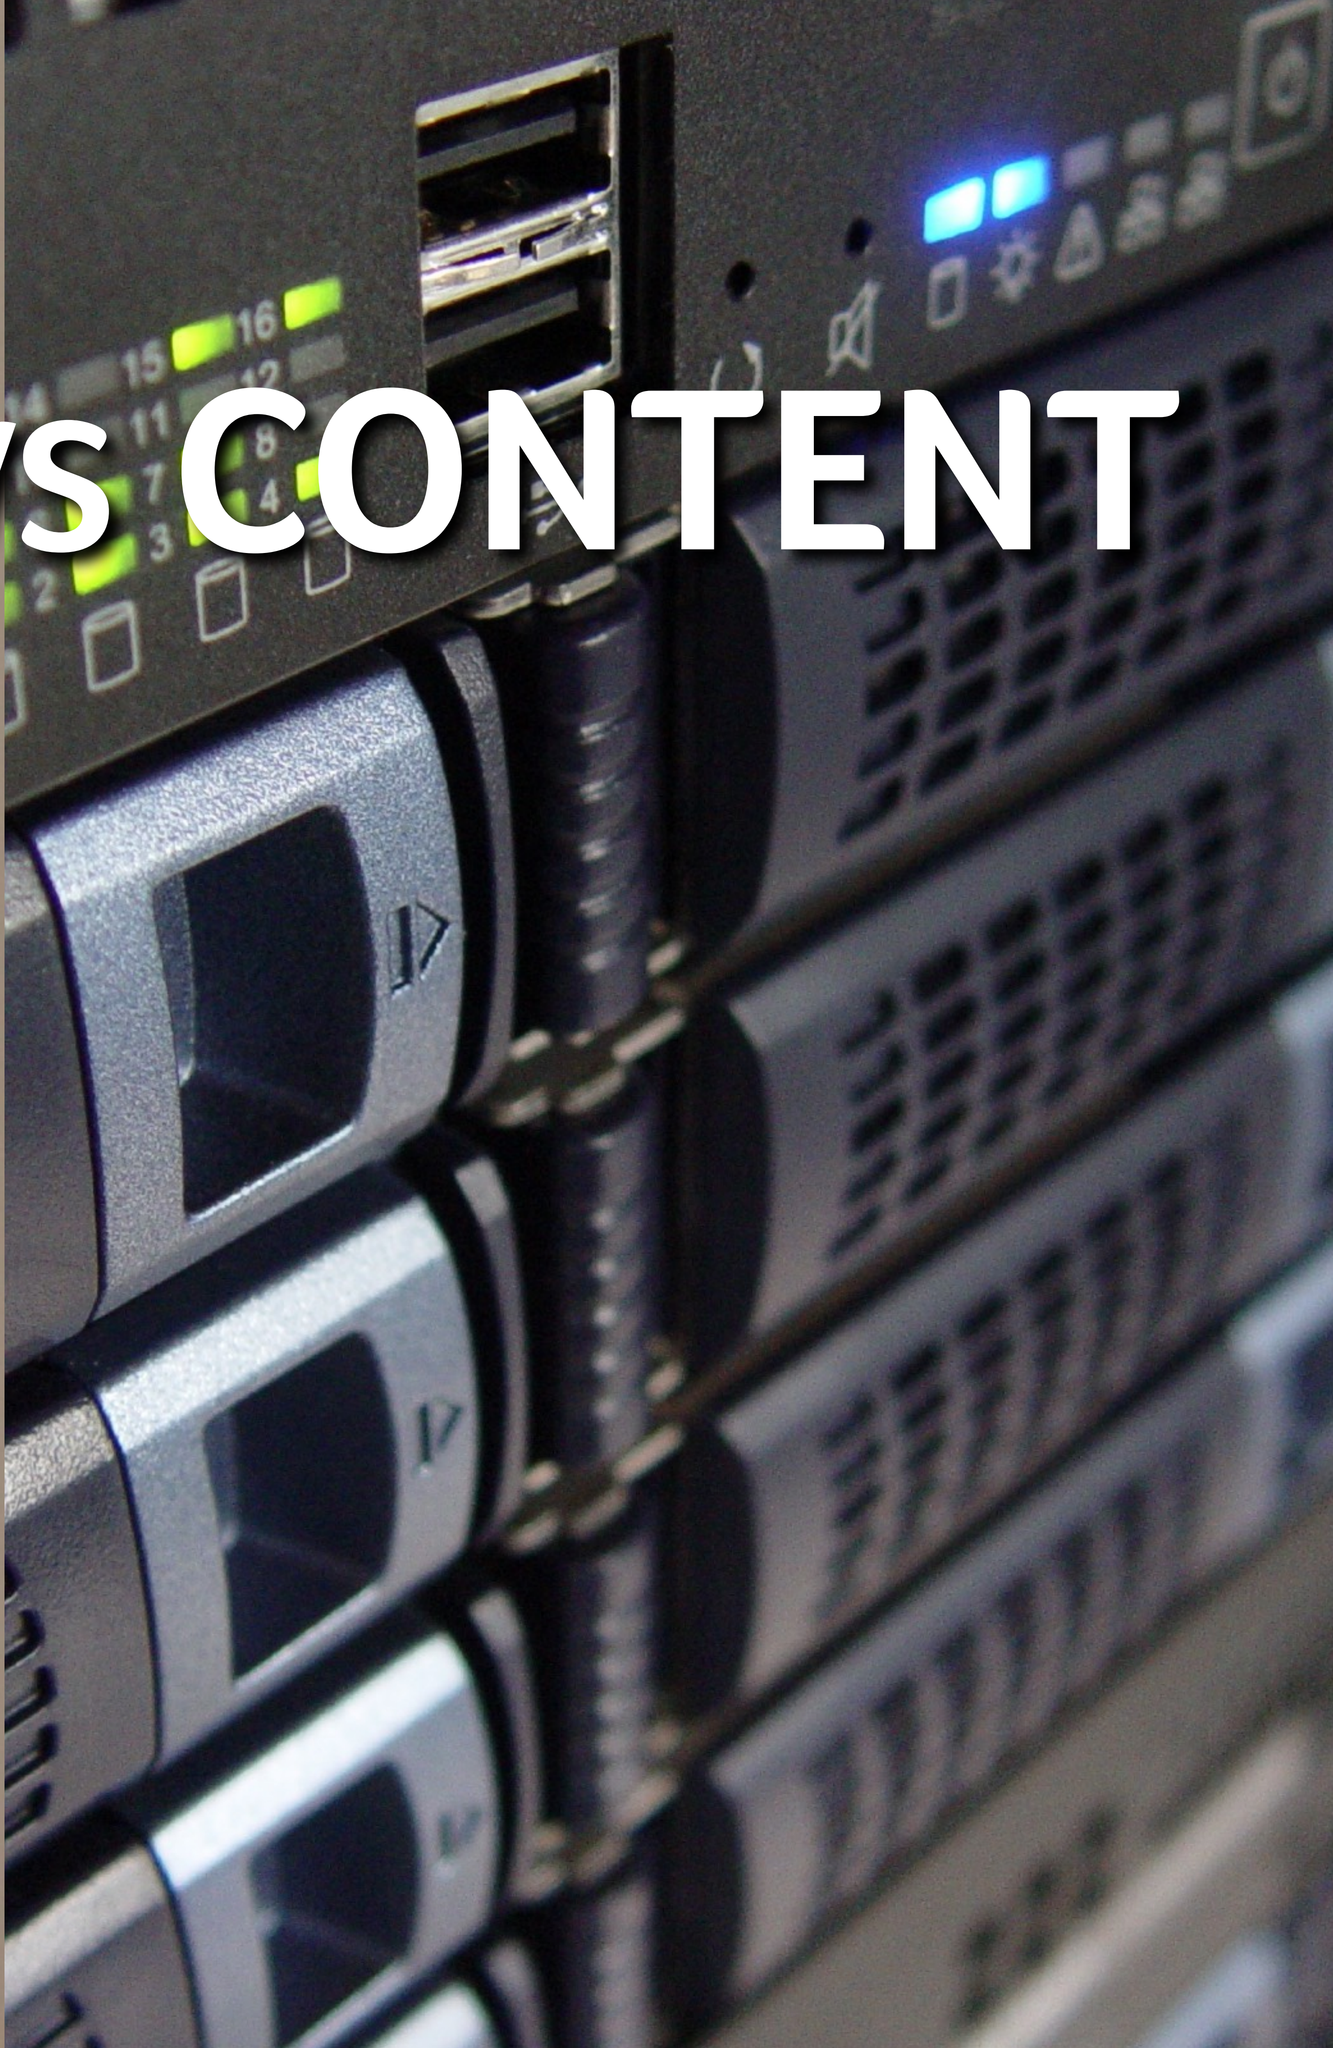

SINGLE vs MULTI

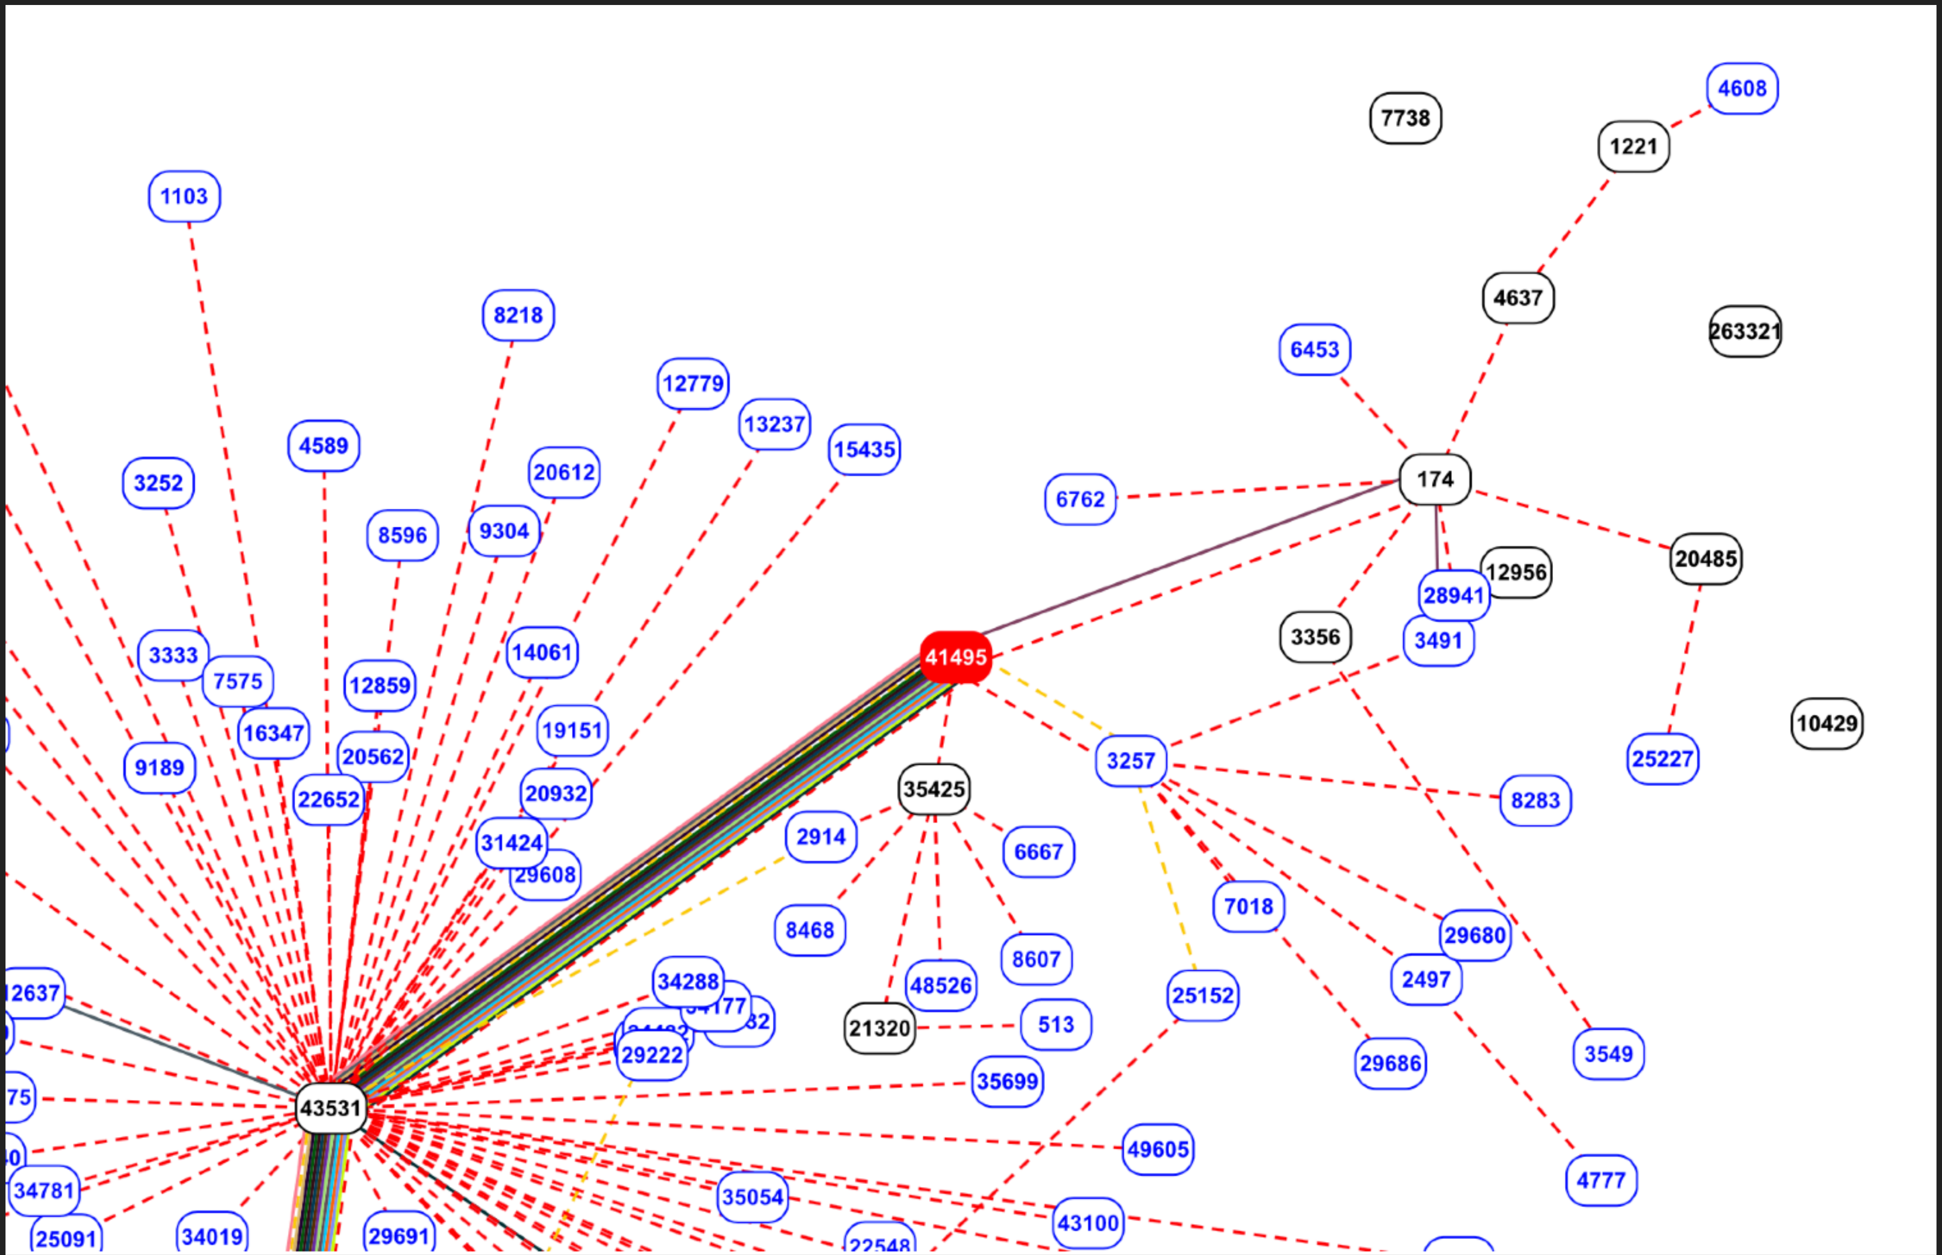

MULTI-HOMED

- Organise "transit" from upstream providers
- Talk BGP with them, announcements + get sent routing tables
- Maybe you get "default only"...

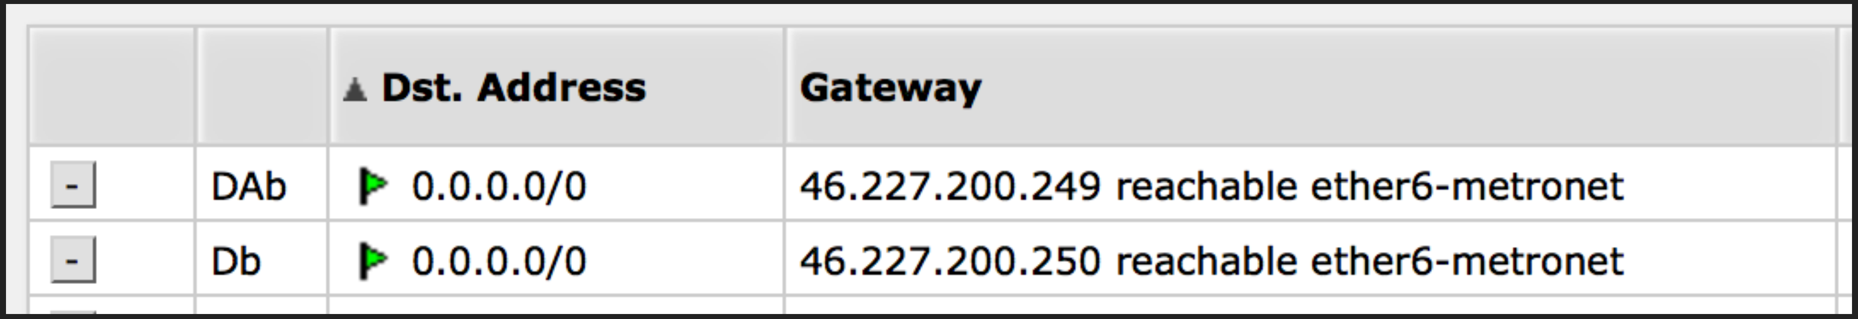

		▲ Dst. Address	Gateway
-	DAb	0.0.0.0/0	46.227.200.249 reachable ether6-metronet
-	Db	0.0.0.0/0	46.227.200.250 reachable ether6-metronet

MULTI-HOMED

- Organise "transit" from upstream providers
- Talk BGP with them, announcements + get sent routing tables
- ...or maybe you get "full tables"
  - >600k IPv4 routes, >30k IPv6 routes
  - That's a lot of routes!

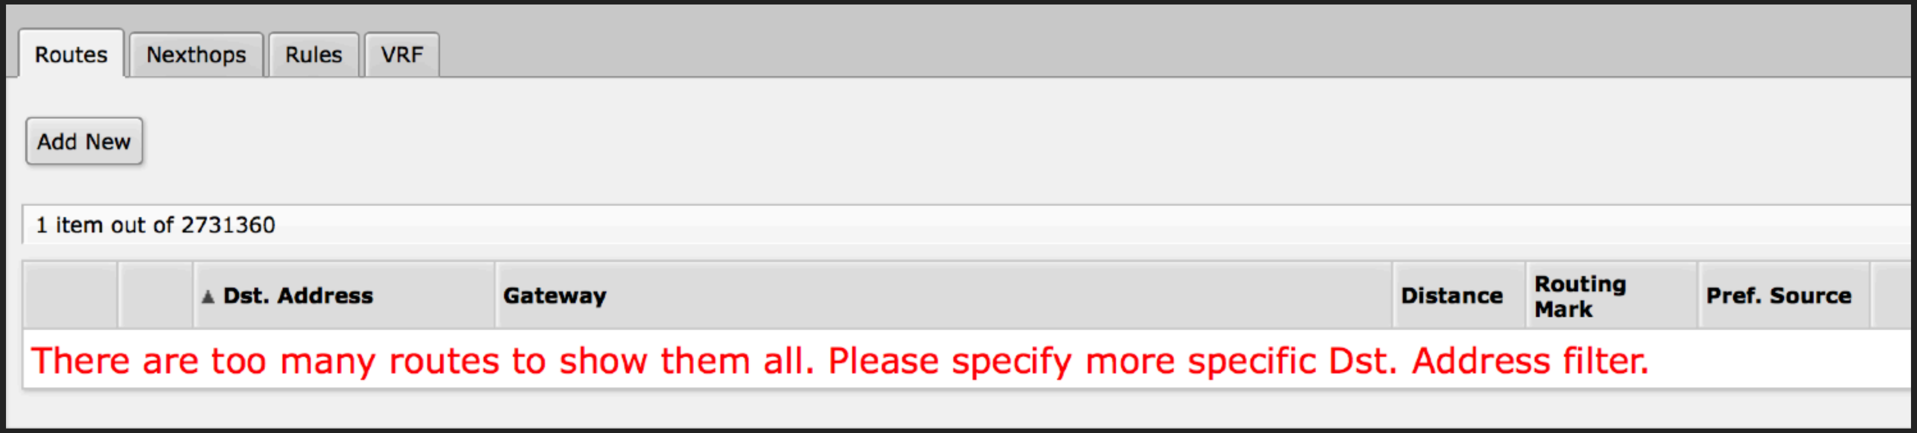

1 item out of 2731360

		▲ Dst. Address	Gateway	Distance	Routing Mark	Pref. Source
There are too many routes to show them all. Please specify more specific Dst. Address filter.						

OUR MIGRATION TO MIKROTIK ROUTEROS

- Quagga + BIRD on servers running Linux solid for >6 years ❤️
- 2015: we wanted to do an upgrade... 📈
- We love the energy efficiency of MikroTik CCR... 💚
- No "NSA/GCHQ inside"... 😍
- Can we use RouterOS? 🤔
  - + BIRD on servers running Linux? 🐧

TWO ROUTING SYSTEMS?

- Early version of BIRD segfaulted, withdrew announcements
  - Quagga kept on running, we did not vanish from DFZ
- Are we sure RouterOS BGP is going to cope?
- What is support going to be like? Debugging?

OVERALL EXPERIENCE

- Some weird behaviour occasionally…
- NTP leap second bug = hard crash
- Disable VLAN interface before changing its physical interface or VID
- Support are helpful and fast; anecdotally, as responsive as the "big name" vendors
- Debugging time = get friendly with RouterOS command-line

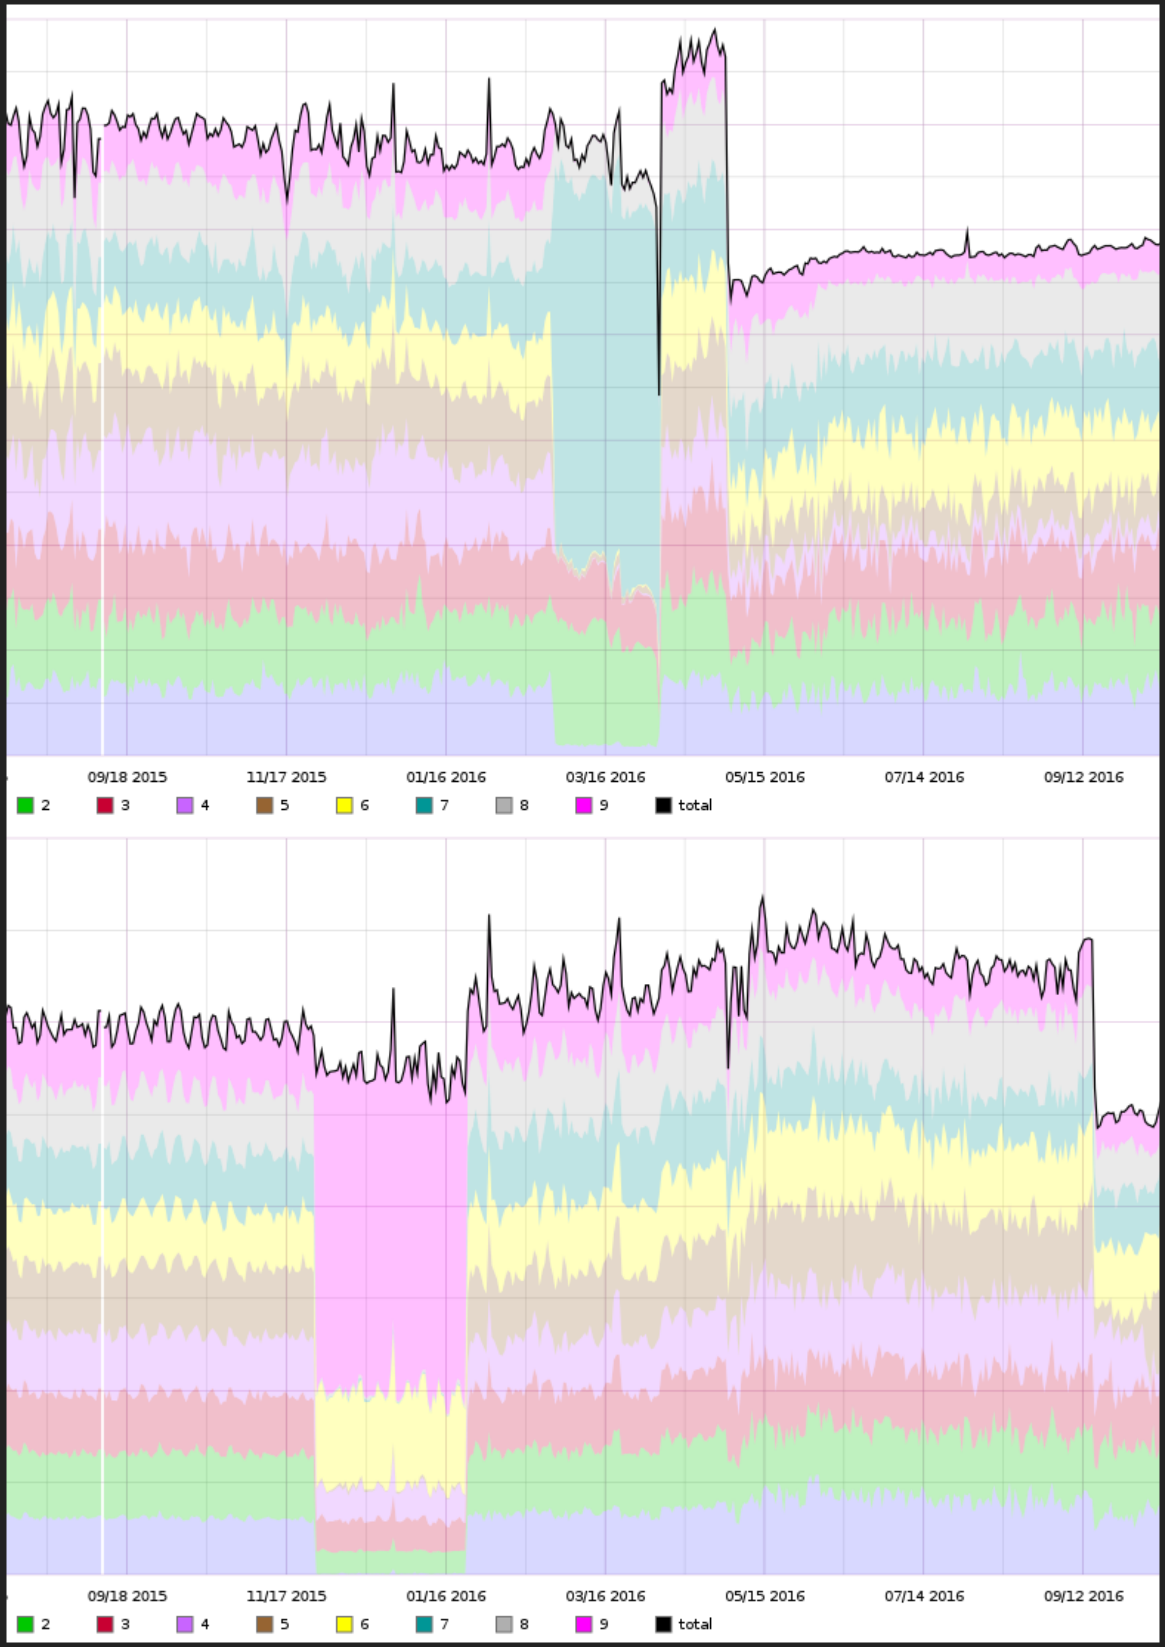

THE GOOD

- £700 + 70W routes >10Gbit/s
- BGP feels familiar after years of experience of Quagga
- Consultants out there if you need them; training & quals
- MikroTik now "go to" choice for CPE, wireless, etc...
- Vendor interop good (beware of extra options in RouterOS)

THE BAD

- Watchdog not good enough, IPMI-style OOB hard reboot?
- BGP converge & FIB is slow on CCR with 2M+ routes
- Routing filters don't always work first time (enable/disable)
- Switch VLAN setup feels like raw config of merchant silicon
- "RouterOS 7"

FAELIX'S TIPS

- CHR, hardware is economical = no excuses for network lab
- Consider leap-frogging RouterOS releases in production
- layer-3 > layer-2, MikroTik affordability = dream come true
- Full routing tables get into FIB a lot quicker on x86 than on tile
- **oxidized** + **syslog** = configs in git + logs in one place
- **snmp** + **graphite** + **grafana** = netops visibility, cool dashboards
- **BCP38** + **MANRS** + **abuse-c** = be excellent to each other

PLUGS

- http://**uknof.org.uk**/ = packet pushers of the UK (and beer)
- http://**netmcr.uk**/ = packet pushers of Manchester (and beer)

PART 2:

FIREWALLING WITH ZERO FILTER RULES!

U WOT M8?

you, right now

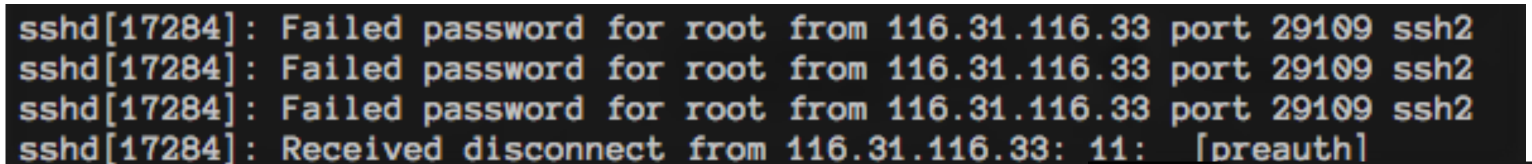

```
sshd[17284]: Failed password for root from 116.31.116.33 port 29109 ssh2
sshd[17284]: Failed password for root from 116.31.116.33 port 29109 ssh2
sshd[17284]: Failed password for root from 116.31.116.33 port 29109 ssh2
sshd[17284]: Received disconnect from 116.31.116.33: 11:  [preauth]
```

ssh

SMTP / IMAP / POP

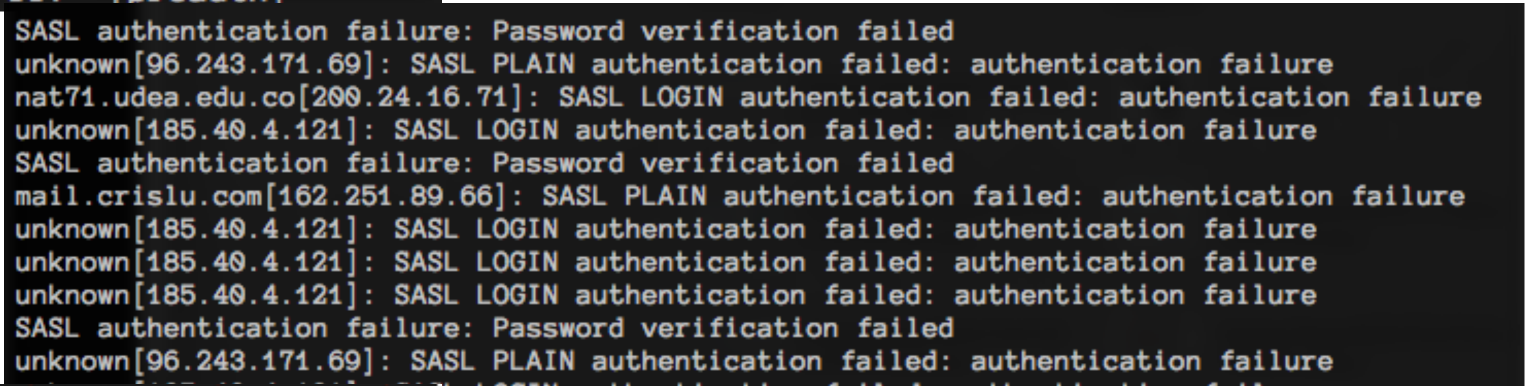

```
SASL authentication failure: Password verification failed
unknown[96.243.171.69]: SASL PLAIN authentication failed: authentication failure
nat71.udea.edu.co[200.24.16.71]: SASL LOGIN authentication failed: authentication failure
unknown[185.40.4.121]: SASL LOGIN authentication failed: authentication failure
SASL authentication failure: Password verification failed
mail.crislu.com[162.251.89.66]: SASL PLAIN authentication failed: authentication failure
unknown[185.40.4.121]: SASL LOGIN authentication failed: authentication failure
unknown[185.40.4.121]: SASL LOGIN authentication failed: authentication failure
unknown[185.40.4.121]: SASL LOGIN authentication failed: authentication failure
SASL authentication failure: Password verification failed
unknown[96.243.171.69]: SASL PLAIN authentication failed: authentication failure
                                L LOGIN authentication failed: authentication failure
```

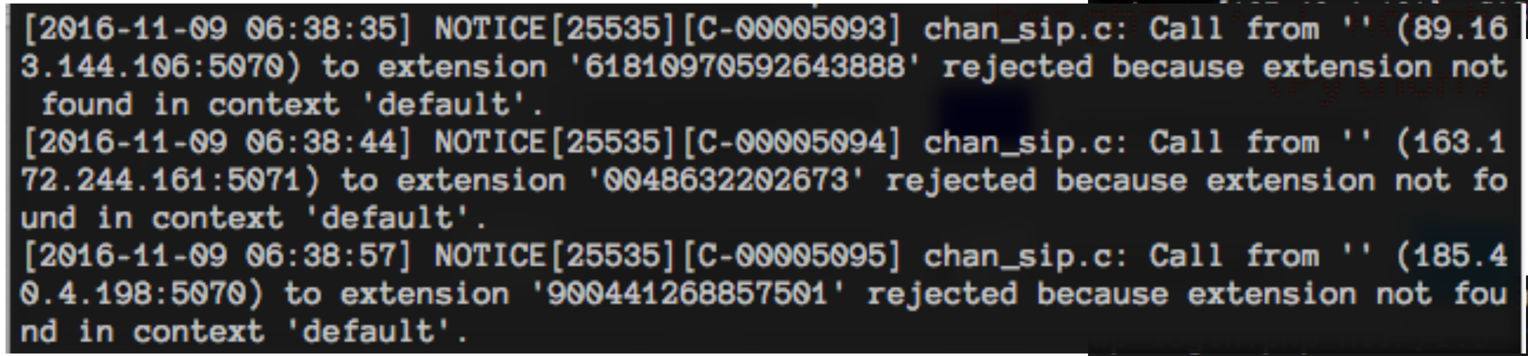

```
[2016-11-09 06:38:35] NOTICE[25535][C-00005093] chan_sip.c: Call from '' (89.16
3.144.106:5070) to extension '61810970592643888' rejected because extension not
 found in context 'default'.
[2016-11-09 06:38:44] NOTICE[25535][C-00005094] chan_sip.c: Call from '' (163.1
72.244.161:5071) to extension '0048632202673' rejected because extension not fo
und in context 'default'.
[2016-11-09 06:38:57] NOTICE[25535][C-00005095] chan_sip.c: Call from '' (185.4
0.4.198:5070) to extension '900441268857501' rejected because extension not fou
nd in context 'default'.
```

VOIP

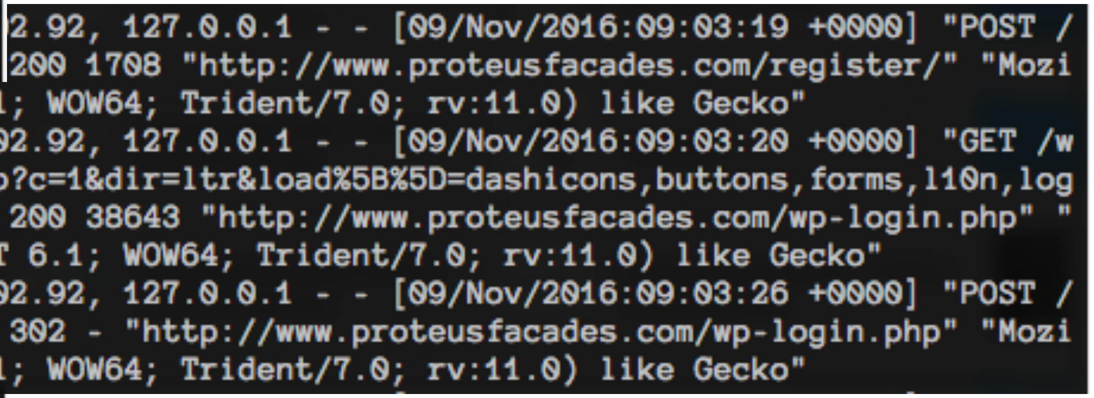

```
2.92, 127.0.0.1 - - [09/Nov/2016:09:03:19 +0000] "POST /
200 1708 "http://www.proteusfacades.com/register/" "Mozi
lla/5.0 (Windows NT 6.1; WOW64; Trident/7.0; rv:11.0) like Gecko"
86.53.243.85, 46.227.202.92, 127.0.0.1 - - [09/Nov/2016:09:03:20 +0000] "GET /w
p-admin/load-styles.php?c=1&dir=ltr&load%5B%5D=dashicons,buttons,forms,l10n,log
in&ver=4.6.1 HTTP/1.0" 200 38643 "http://www.proteusfacades.com/wp-login.php" "
Mozilla/5.0 (Windows NT 6.1; WOW64; Trident/7.0; rv:11.0) like Gecko"
86.53.243.85, 46.227.202.92, 127.0.0.1 - - [09/Nov/2016:09:03:26 +0000] "POST /
wp-login.php HTTP/1.0" 302 - "http://www.proteusfacades.com/wp-login.php" "Mozi
lla/5.0 (Windows NT 6.1; WOW64; Trident/7.0; rv:11.0) like Gecko"
```

WordPress

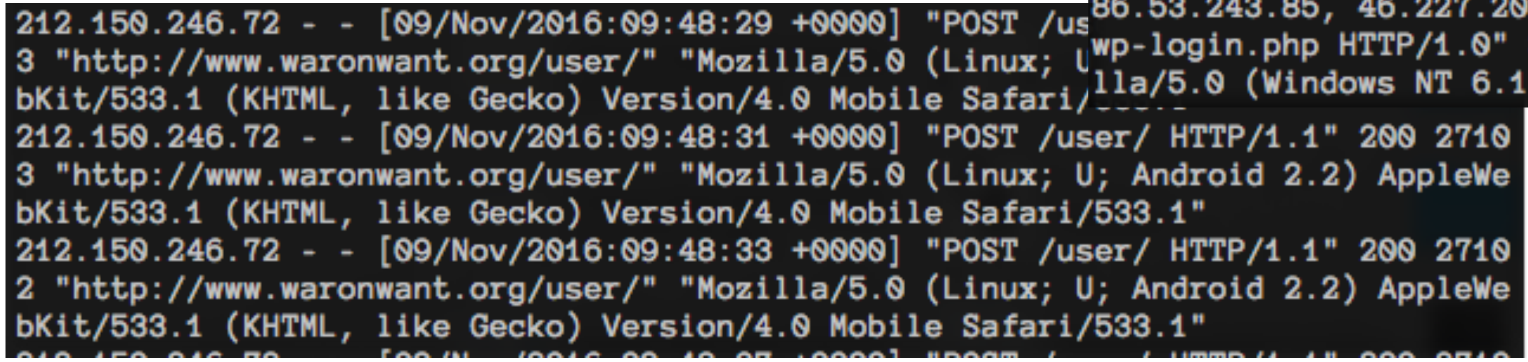

```
212.150.246.72 - - [09/Nov/2016:09:48:29 +0000] "POST /us
3 "http://www.waronwant.org/user/" "Mozilla/5.0 (Linux; U
bKit/533.1 (KHTML, like Gecko) Version/4.0 Mobile Safari/
212.150.246.72 - - [09/Nov/2016:09:48:31 +0000] "POST /user/ HTTP/1.1" 200 2710
3 "http://www.waronwant.org/user/" "Mozilla/5.0 (Linux; U; Android 2.2) AppleWe
bKit/533.1 (KHTML, like Gecko) Version/4.0 Mobile Safari/533.1"
212.150.246.72 - - [09/Nov/2016:09:48:33 +0000] "POST /user/ HTTP/1.1" 200 2710
2 "http://www.waronwant.org/user/" "Mozilla/5.0 (Linux; U; Android 2.2) AppleWe
bKit/533.1 (KHTML, like Gecko) Version/4.0 Mobile Safari/533.1"
```

Drupal

L2TP

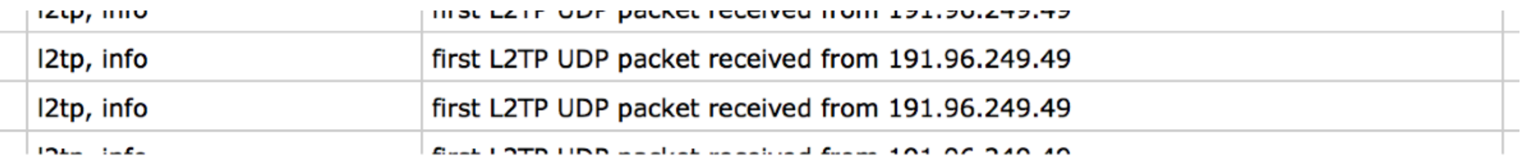

l2tp, info	first L2TP UDP packet received from 191.96.249.49
l2tp, info	first L2TP UDP packet received from 191.96.249.49

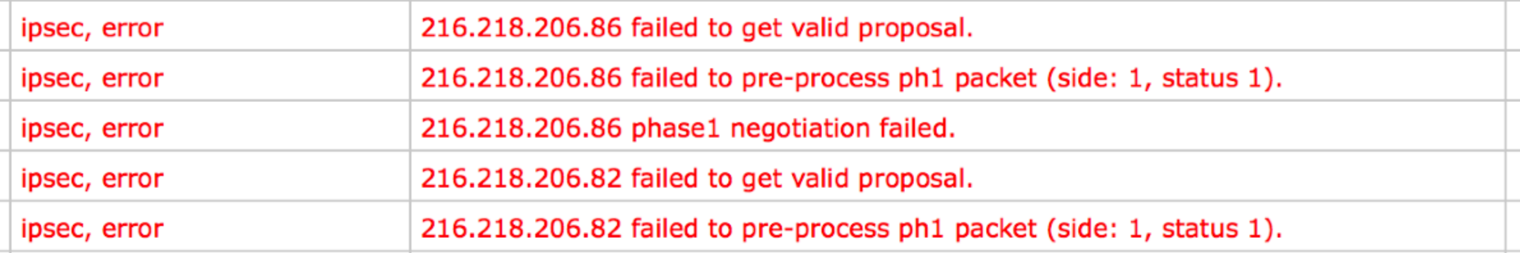

ipsec, error	216.218.206.86 failed to get valid proposal.
ipsec, error	216.218.206.86 failed to pre-process ph1 packet (side: 1, status 1).
ipsec, error	216.218.206.86 phase1 negotiation failed.
ipsec, error	216.218.206.82 failed to get valid proposal.
ipsec, error	216.218.206.82 failed to pre-process ph1 packet (side: 1, status 1).

IPsec

SHIT HAPPENS

- Your network will get scanned
  - ssh, DDoS amplification, open proxies...
- You might have forgotten something
  - Is your management network isolated?
- Your customers will do things you don't expect
  - e.g. SNMP or DNS on CPE open to Internet
- Software has bugs

That is one big pile of shit!
MOVIECLIPS.COM

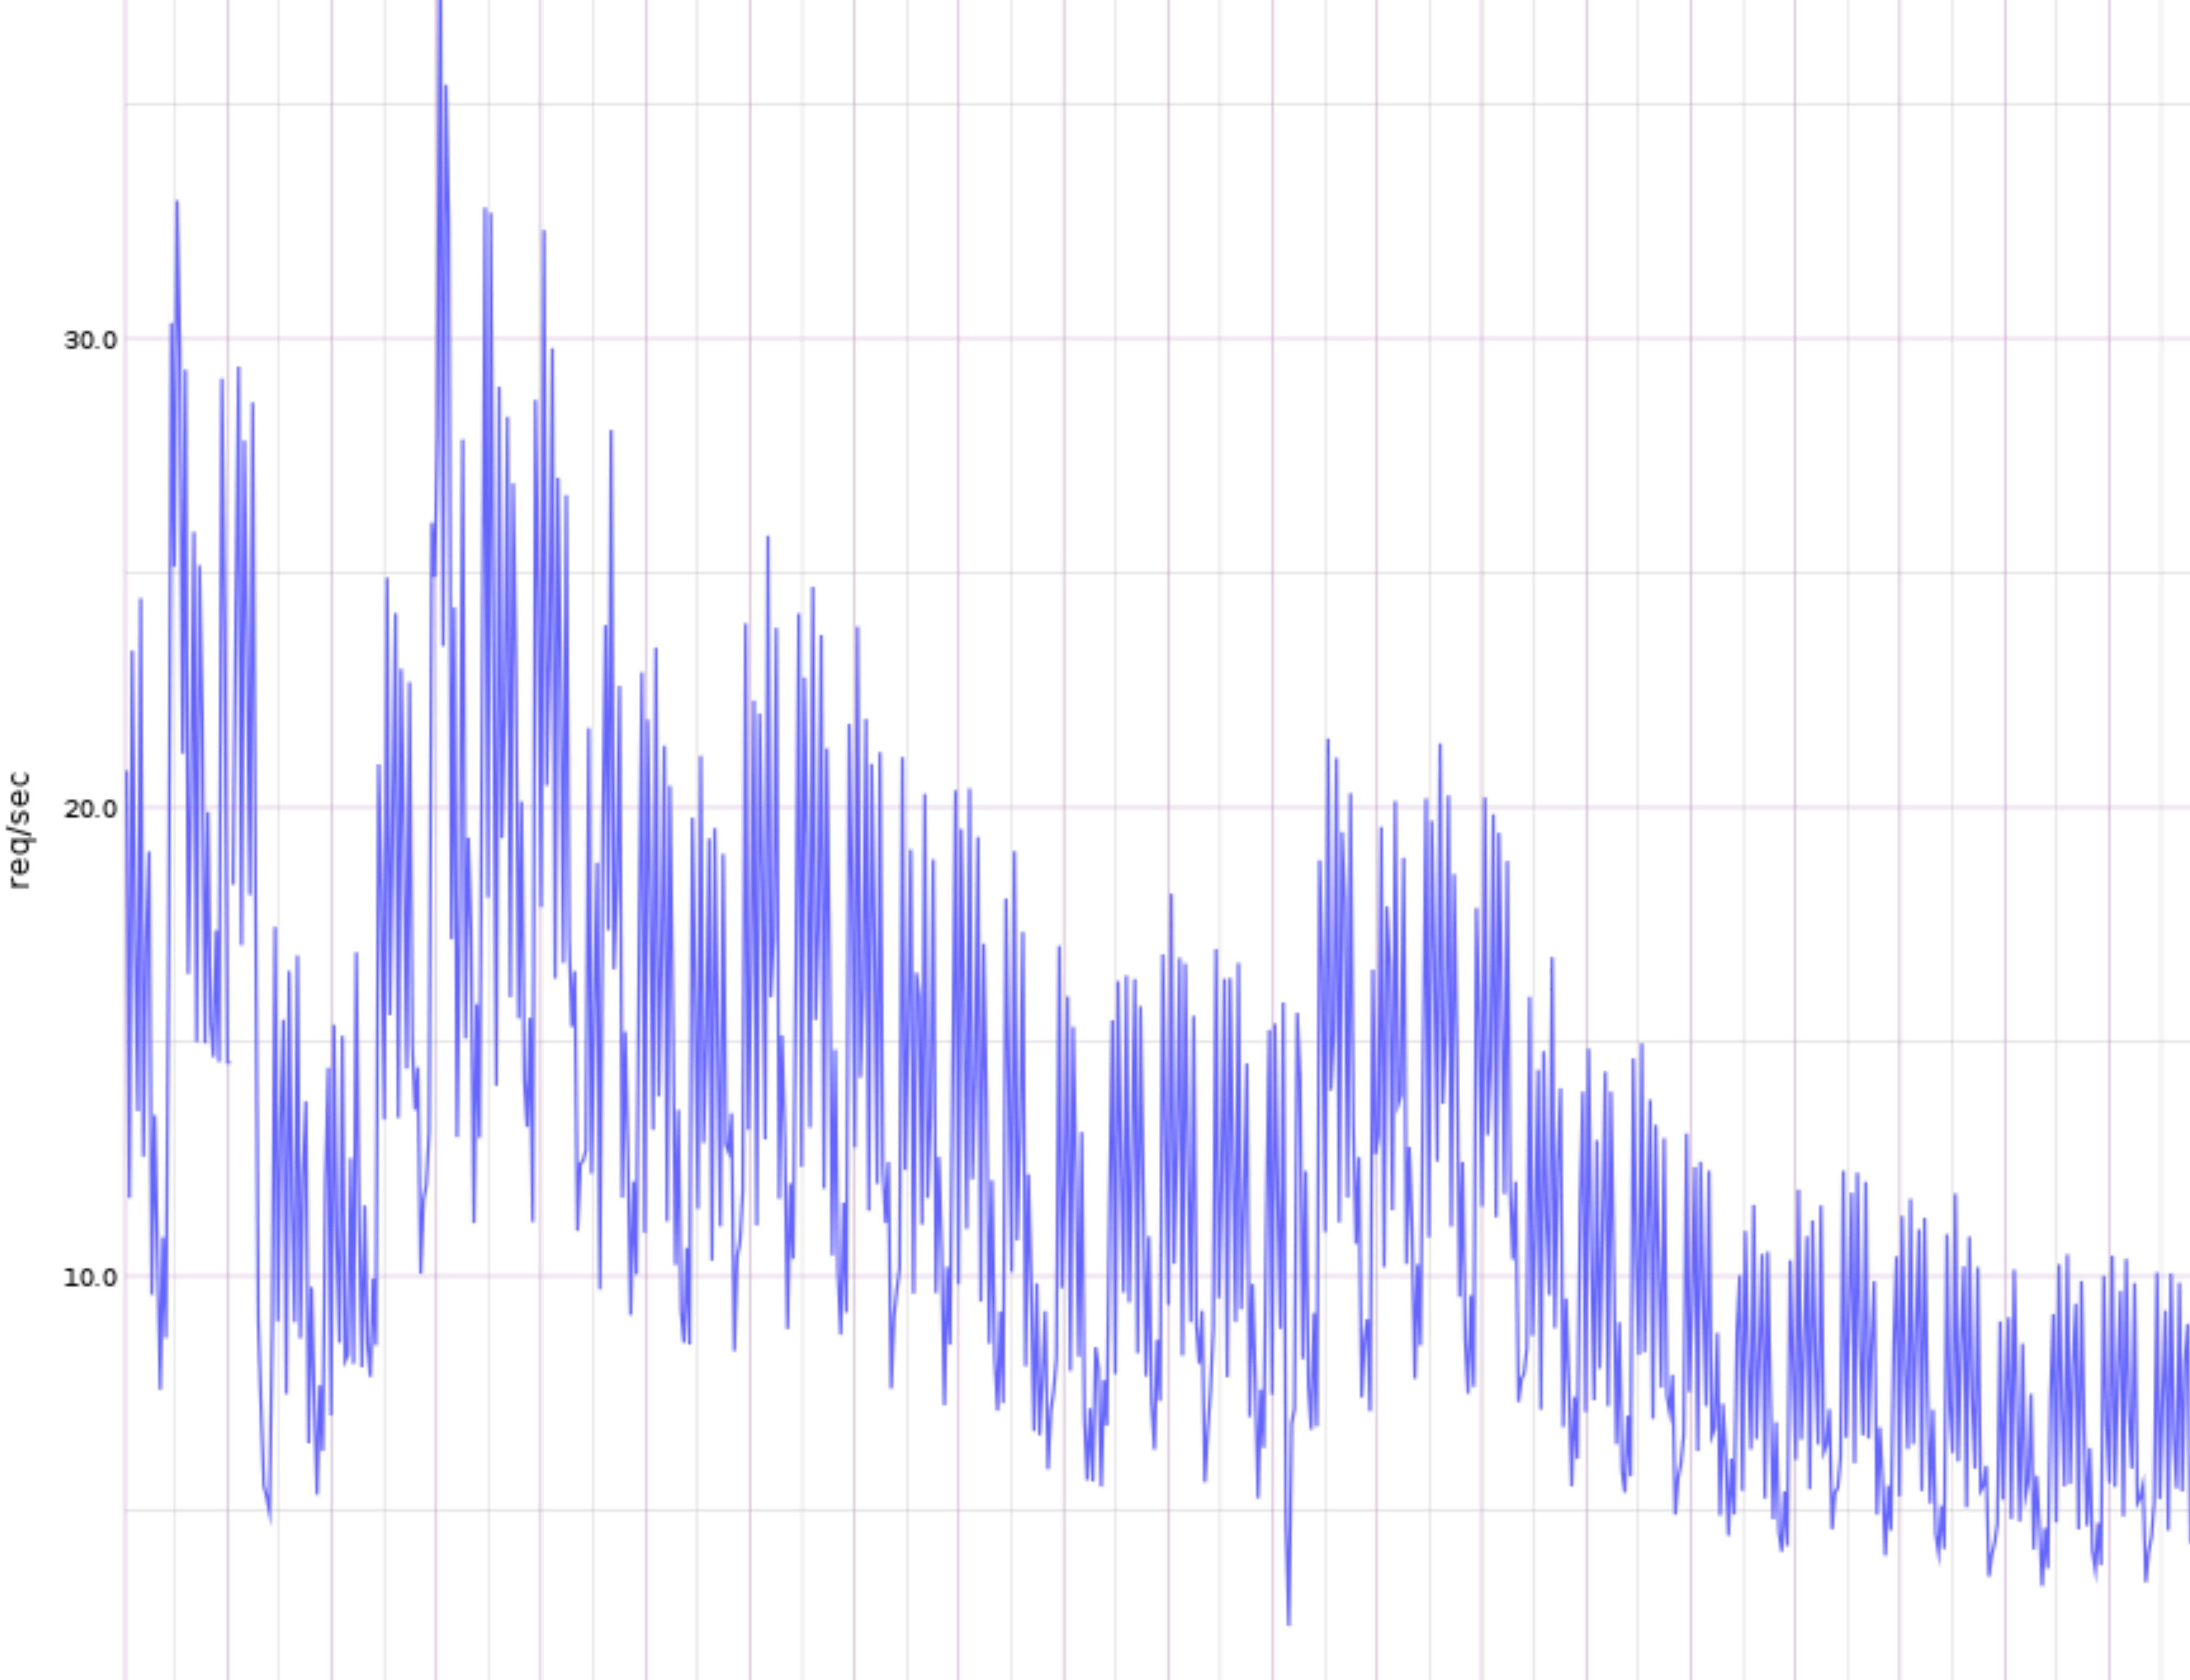
30.0
20.0
10.0
req/sec

omg wtf loadavg bbq

> **SECURITY IS HARD**

every infosec professional ever

GOAL:

START WITH THE LOW-HANGING FRUIT…

AND WHEN THEY'RE PICKED...

THE NEXT CROP!

STEP 1:

LOGS + DATA

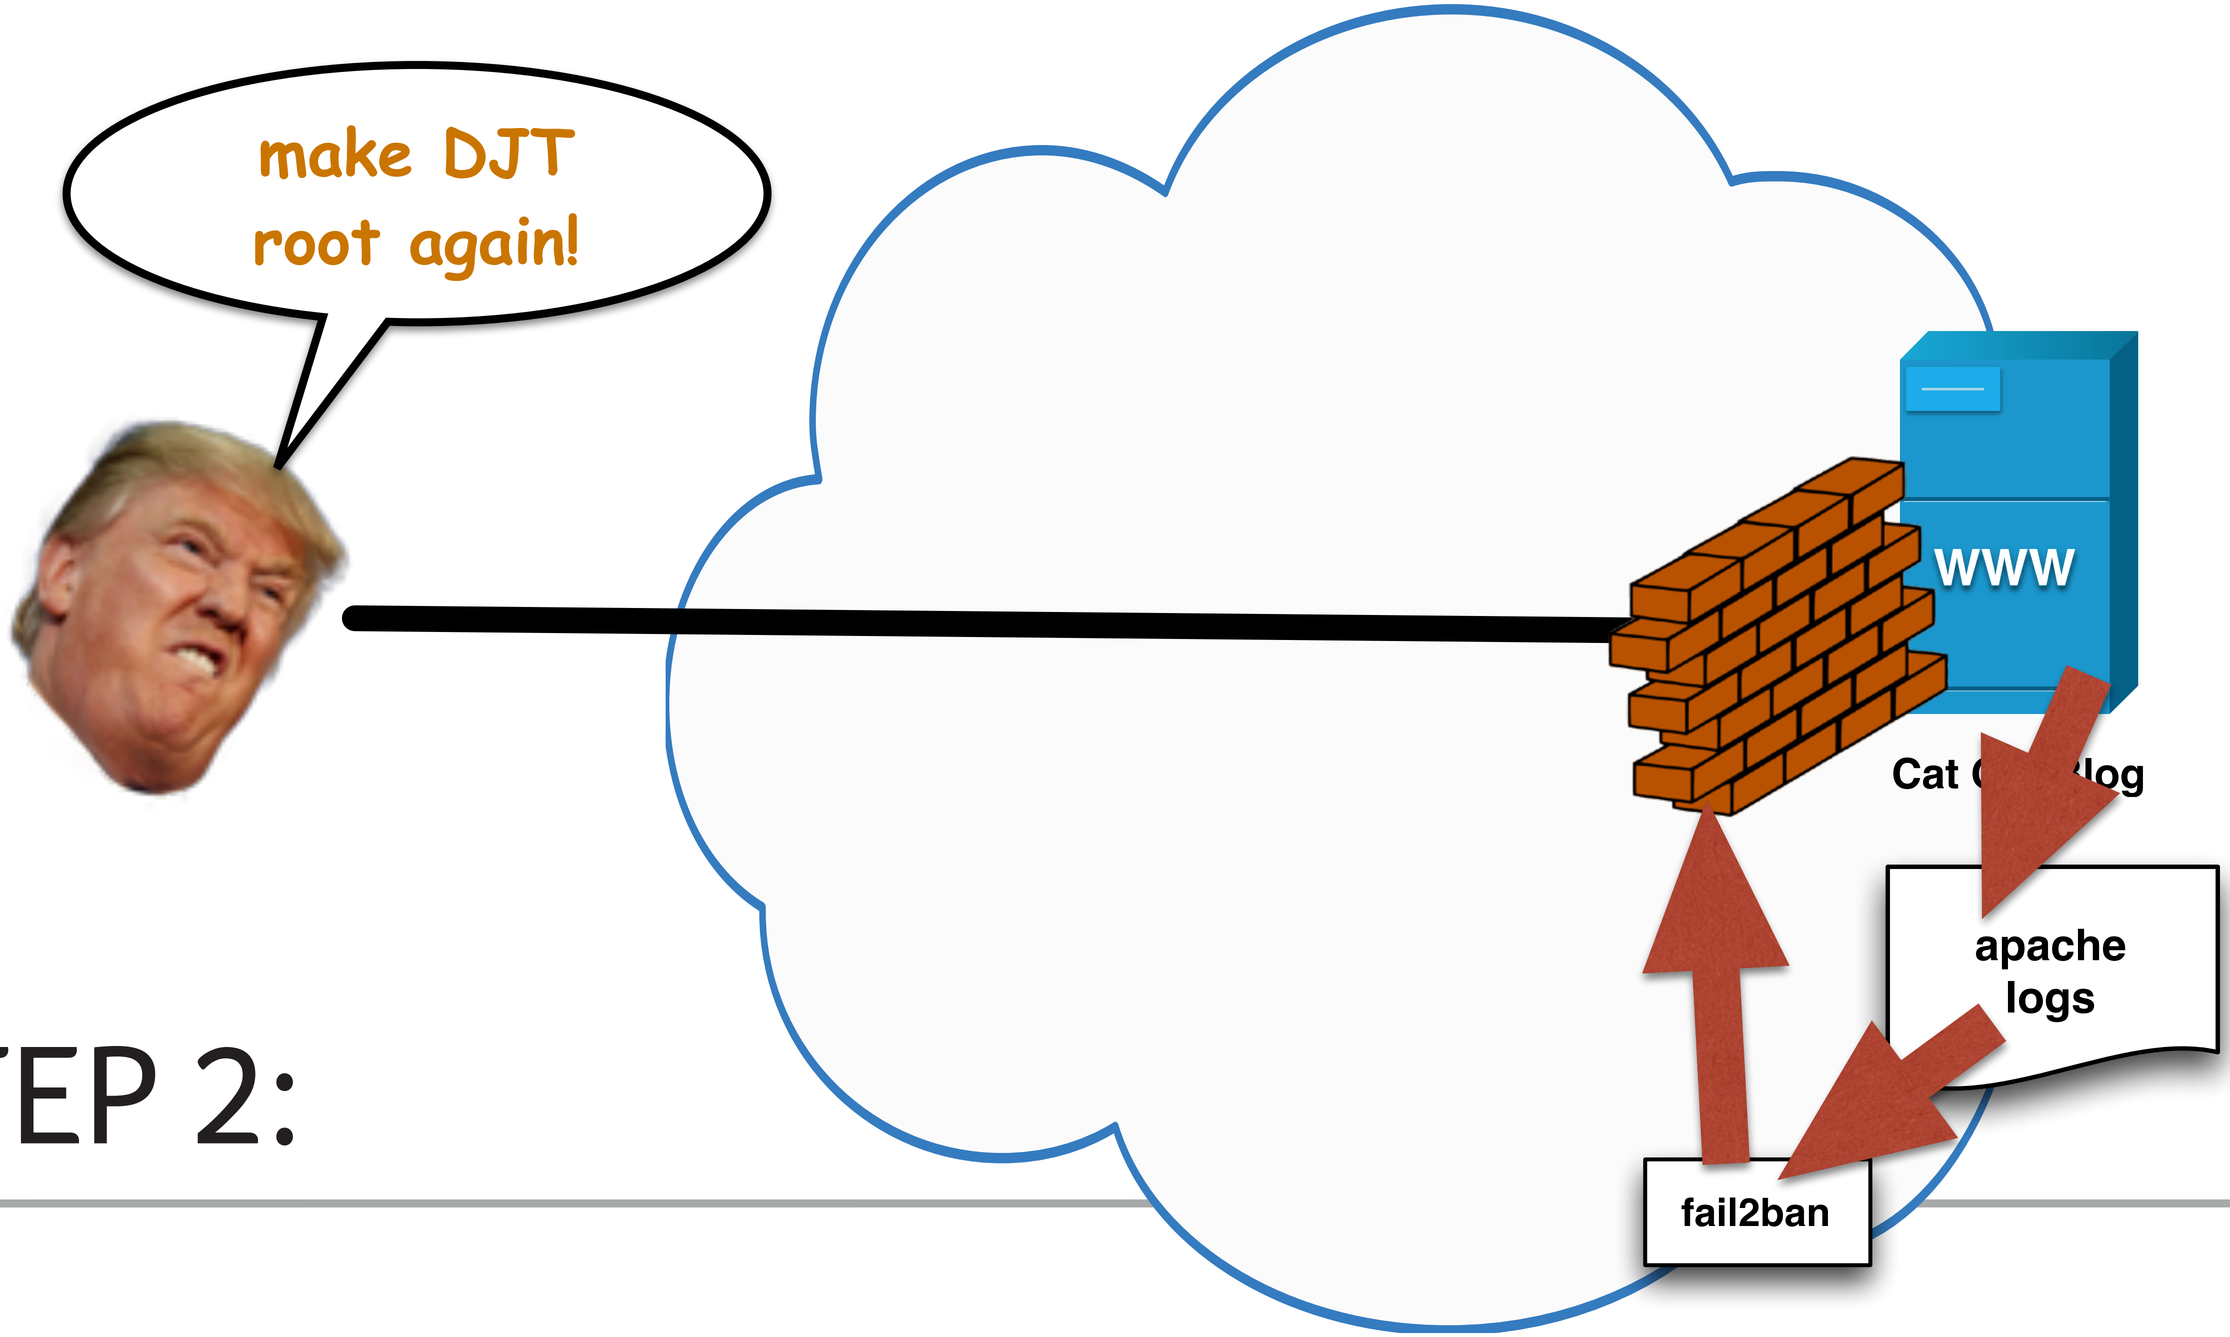

STEP 2:

FAIL2BAN

FAIL2BAN

- Follow log file, if line matches "filter" then performs "action"
- Great for blocking brute force (ssh, etc)
- MikroTik wiki + forum have examples for RouterOS
  - Send logs via syslog to a VM for analysis
  - fail2ban connects to RouterOS with ssh and blocks using:
    - add new **/ip firewall filter** (ok)
    - add new **/ip firewall address-list** (better)

FAIL2BAN

- Quick, cheap, easy
- Make your own or find rules to block web, VoIP, and other nasty traffic
- Attacker will move on to another target pretty quickly when DROPped
- Next target might still be in your network, still traffic across your backbone
- Can we put attacking IPs on a network-wide "naughty step"?

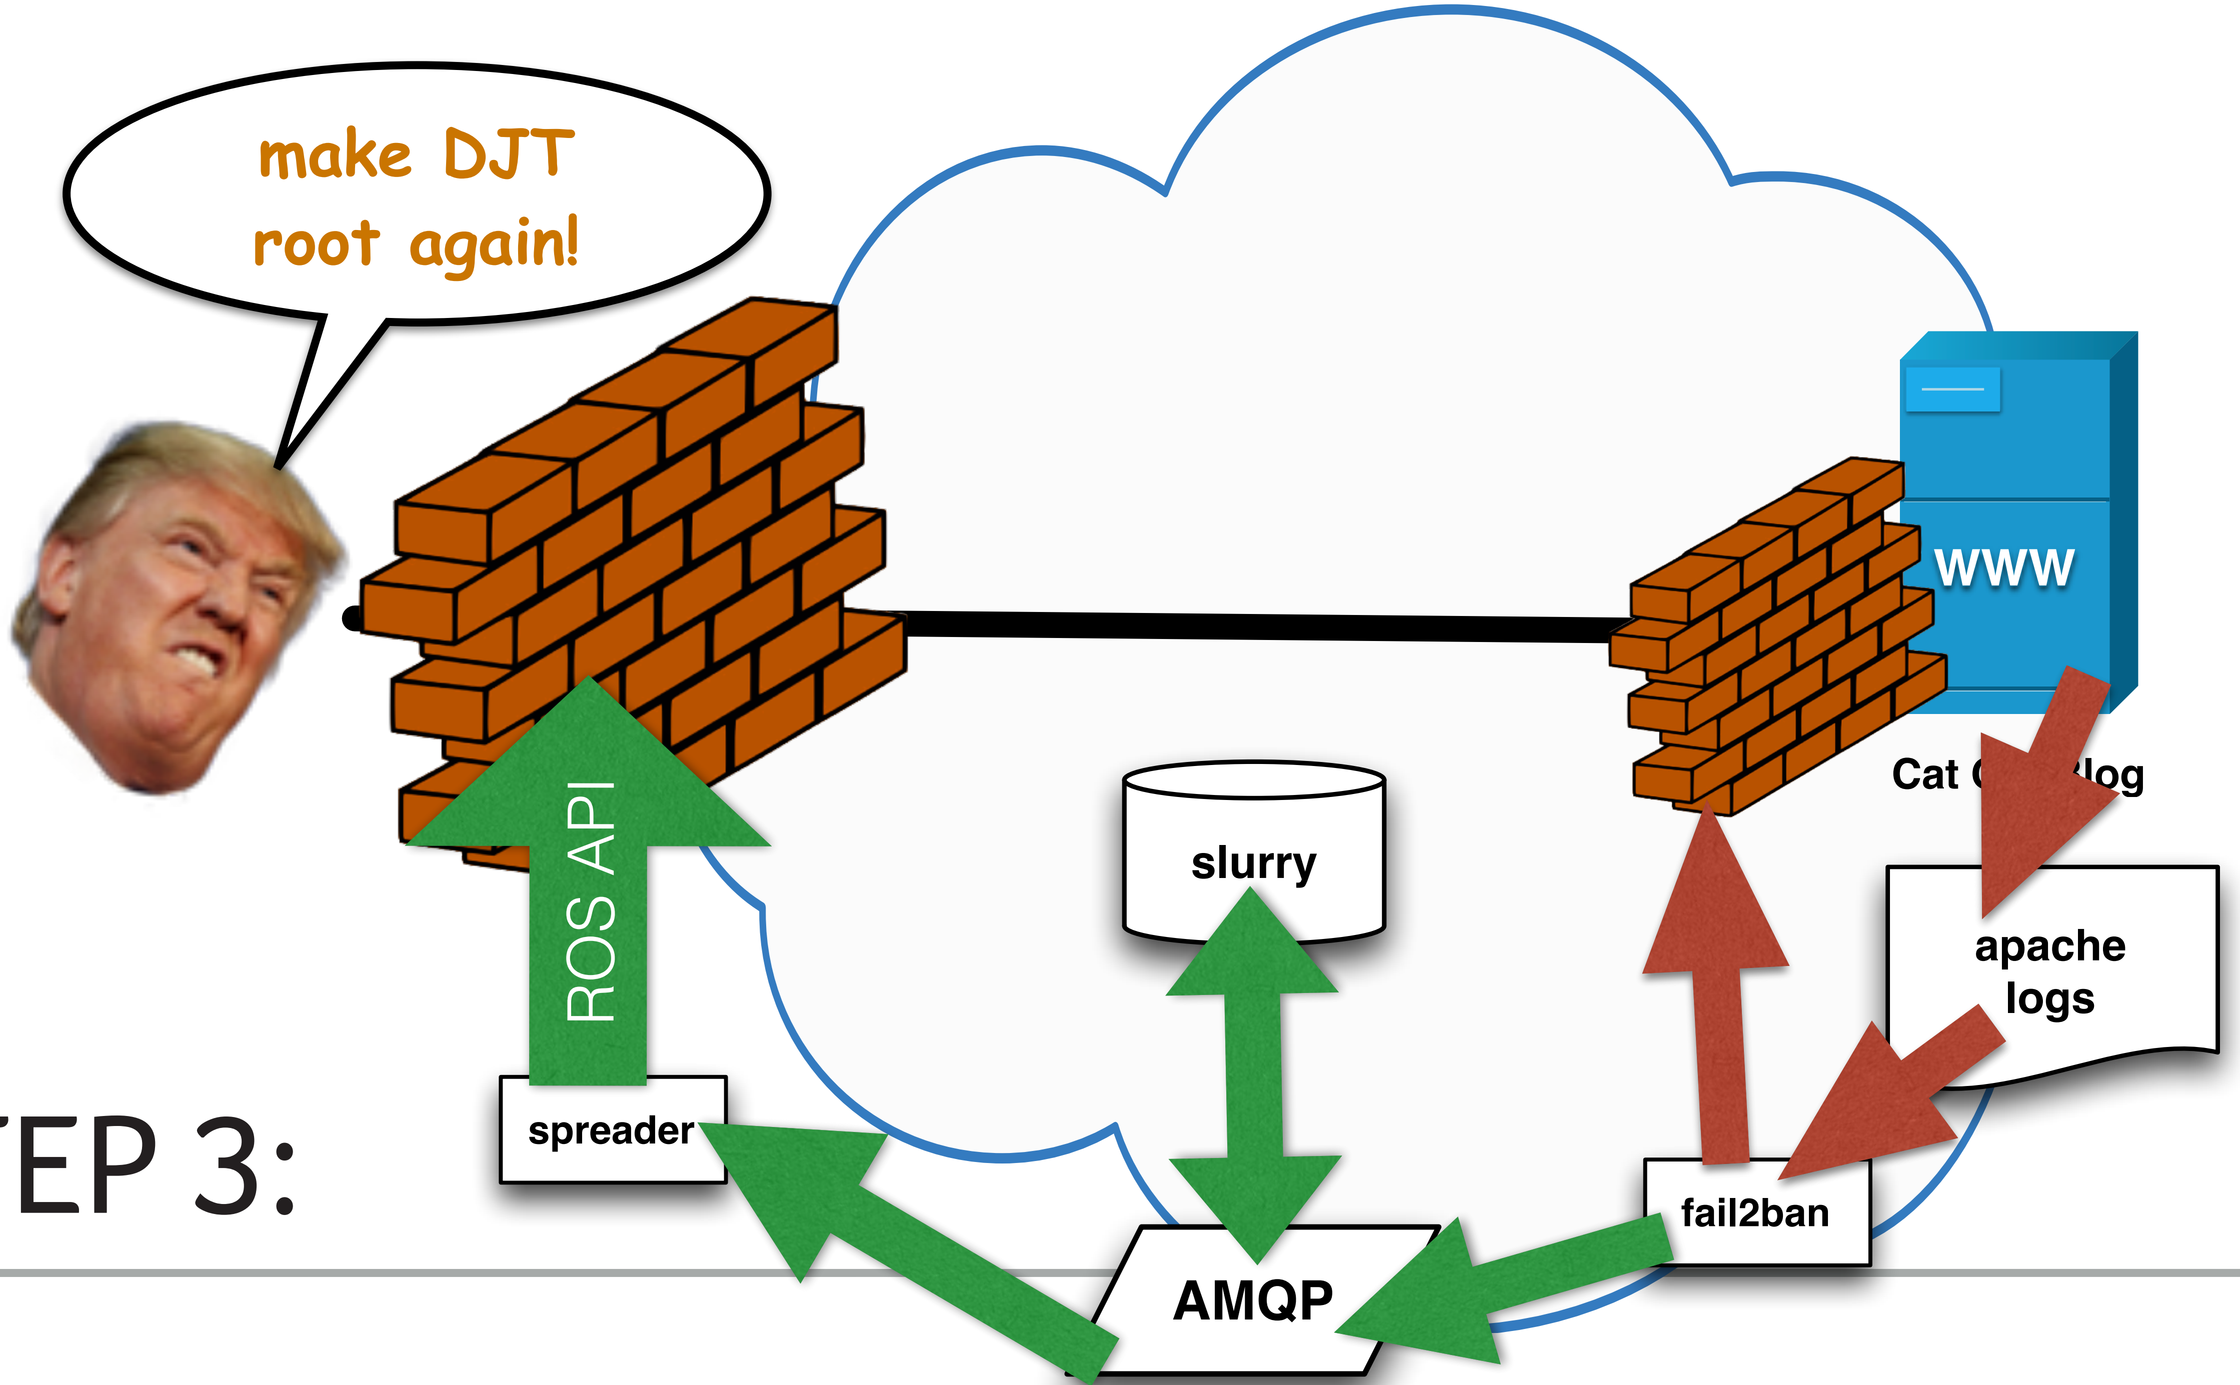

STEP 3:

BLOCK AT THE PROVIDER EDGE

BLOCKING AT THE PROVIDER EDGE

- Lots of flows, lots of PPS, lots of attacking addresses
  - **/ip firewall filter** uses each set of rules sequentially = O(n)
  - **/ip firewall address-list** is a hash-table ≈ O(1)
- Using AMQP to get addresses added to block lists on all routers in three data-centres
  - We already had RabbitMQ across our network for other infrastructure needs

STEP 4:

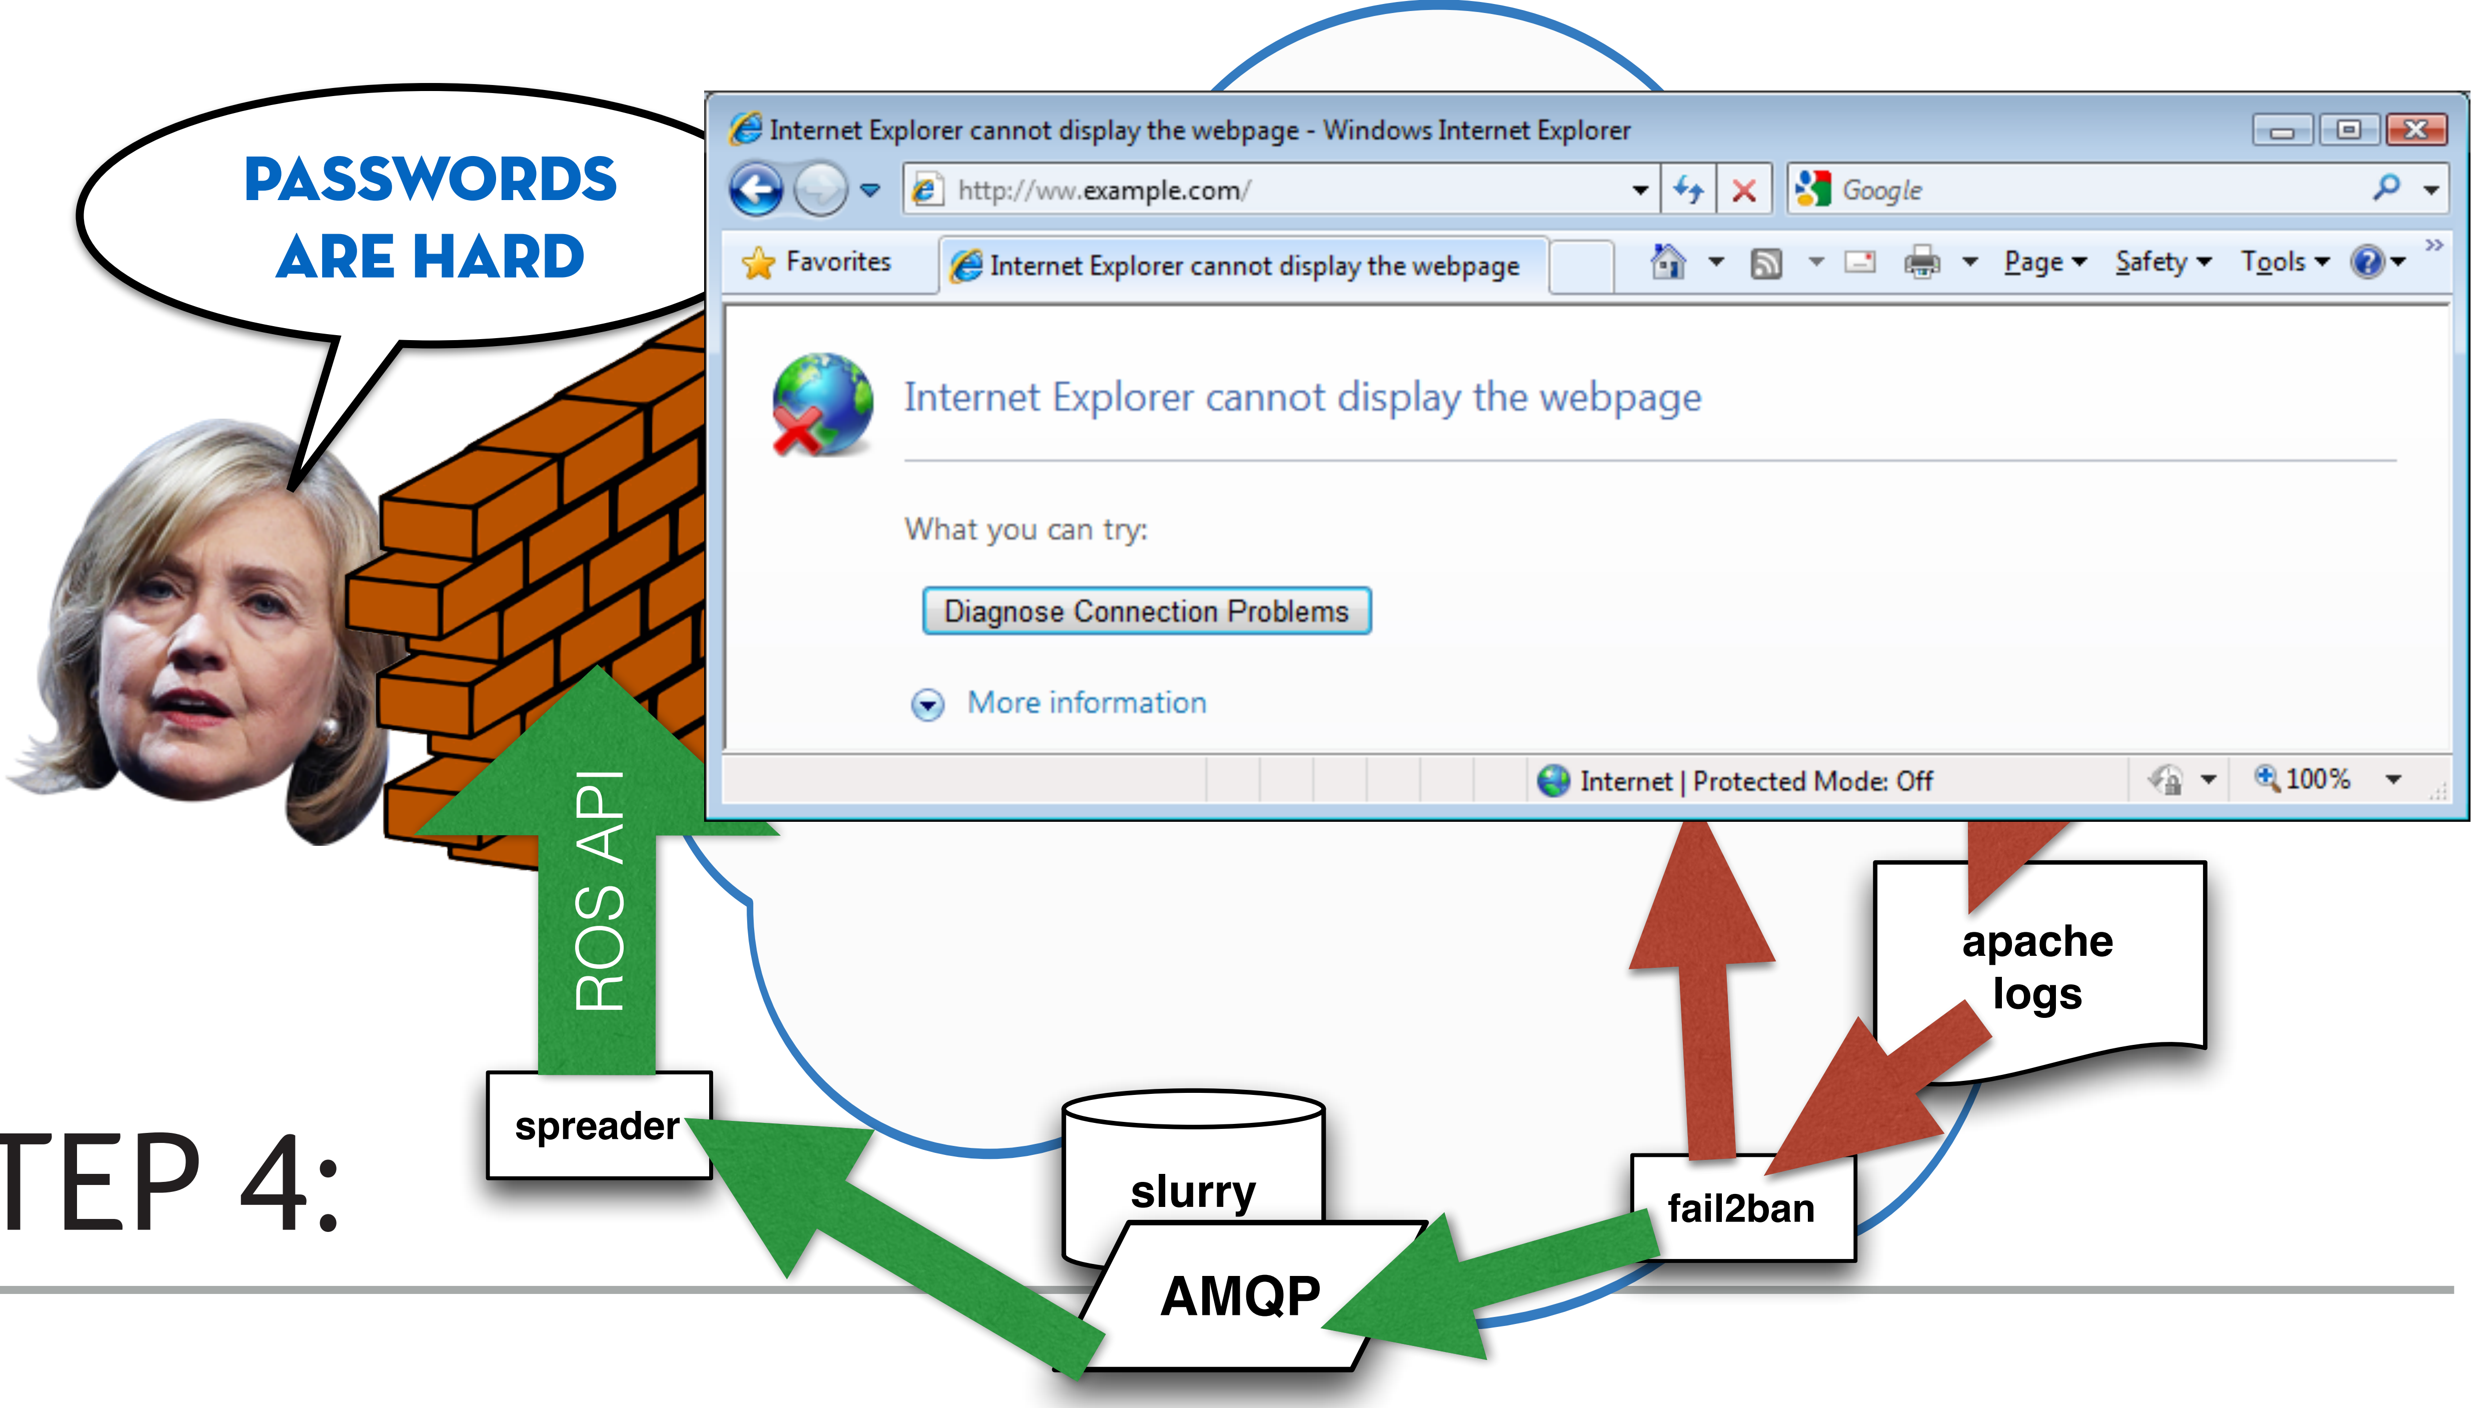
PASSWORDS ARE HARD
ROS API
spreader
slurry
AMQP
fail2ban
apache logs
Internet Explorer cannot display the webpage - Windows Internet Explorer
http://ww.example.com/
Google
Favorites
Internet Explorer cannot display the webpage
Page
Safety
Tools
Diagnose Connection Problems
More information
Internet | Protected Mode: Off
100%

Internet Explorer cannot display the webpage

What you can try:

FALSE POSITIVES

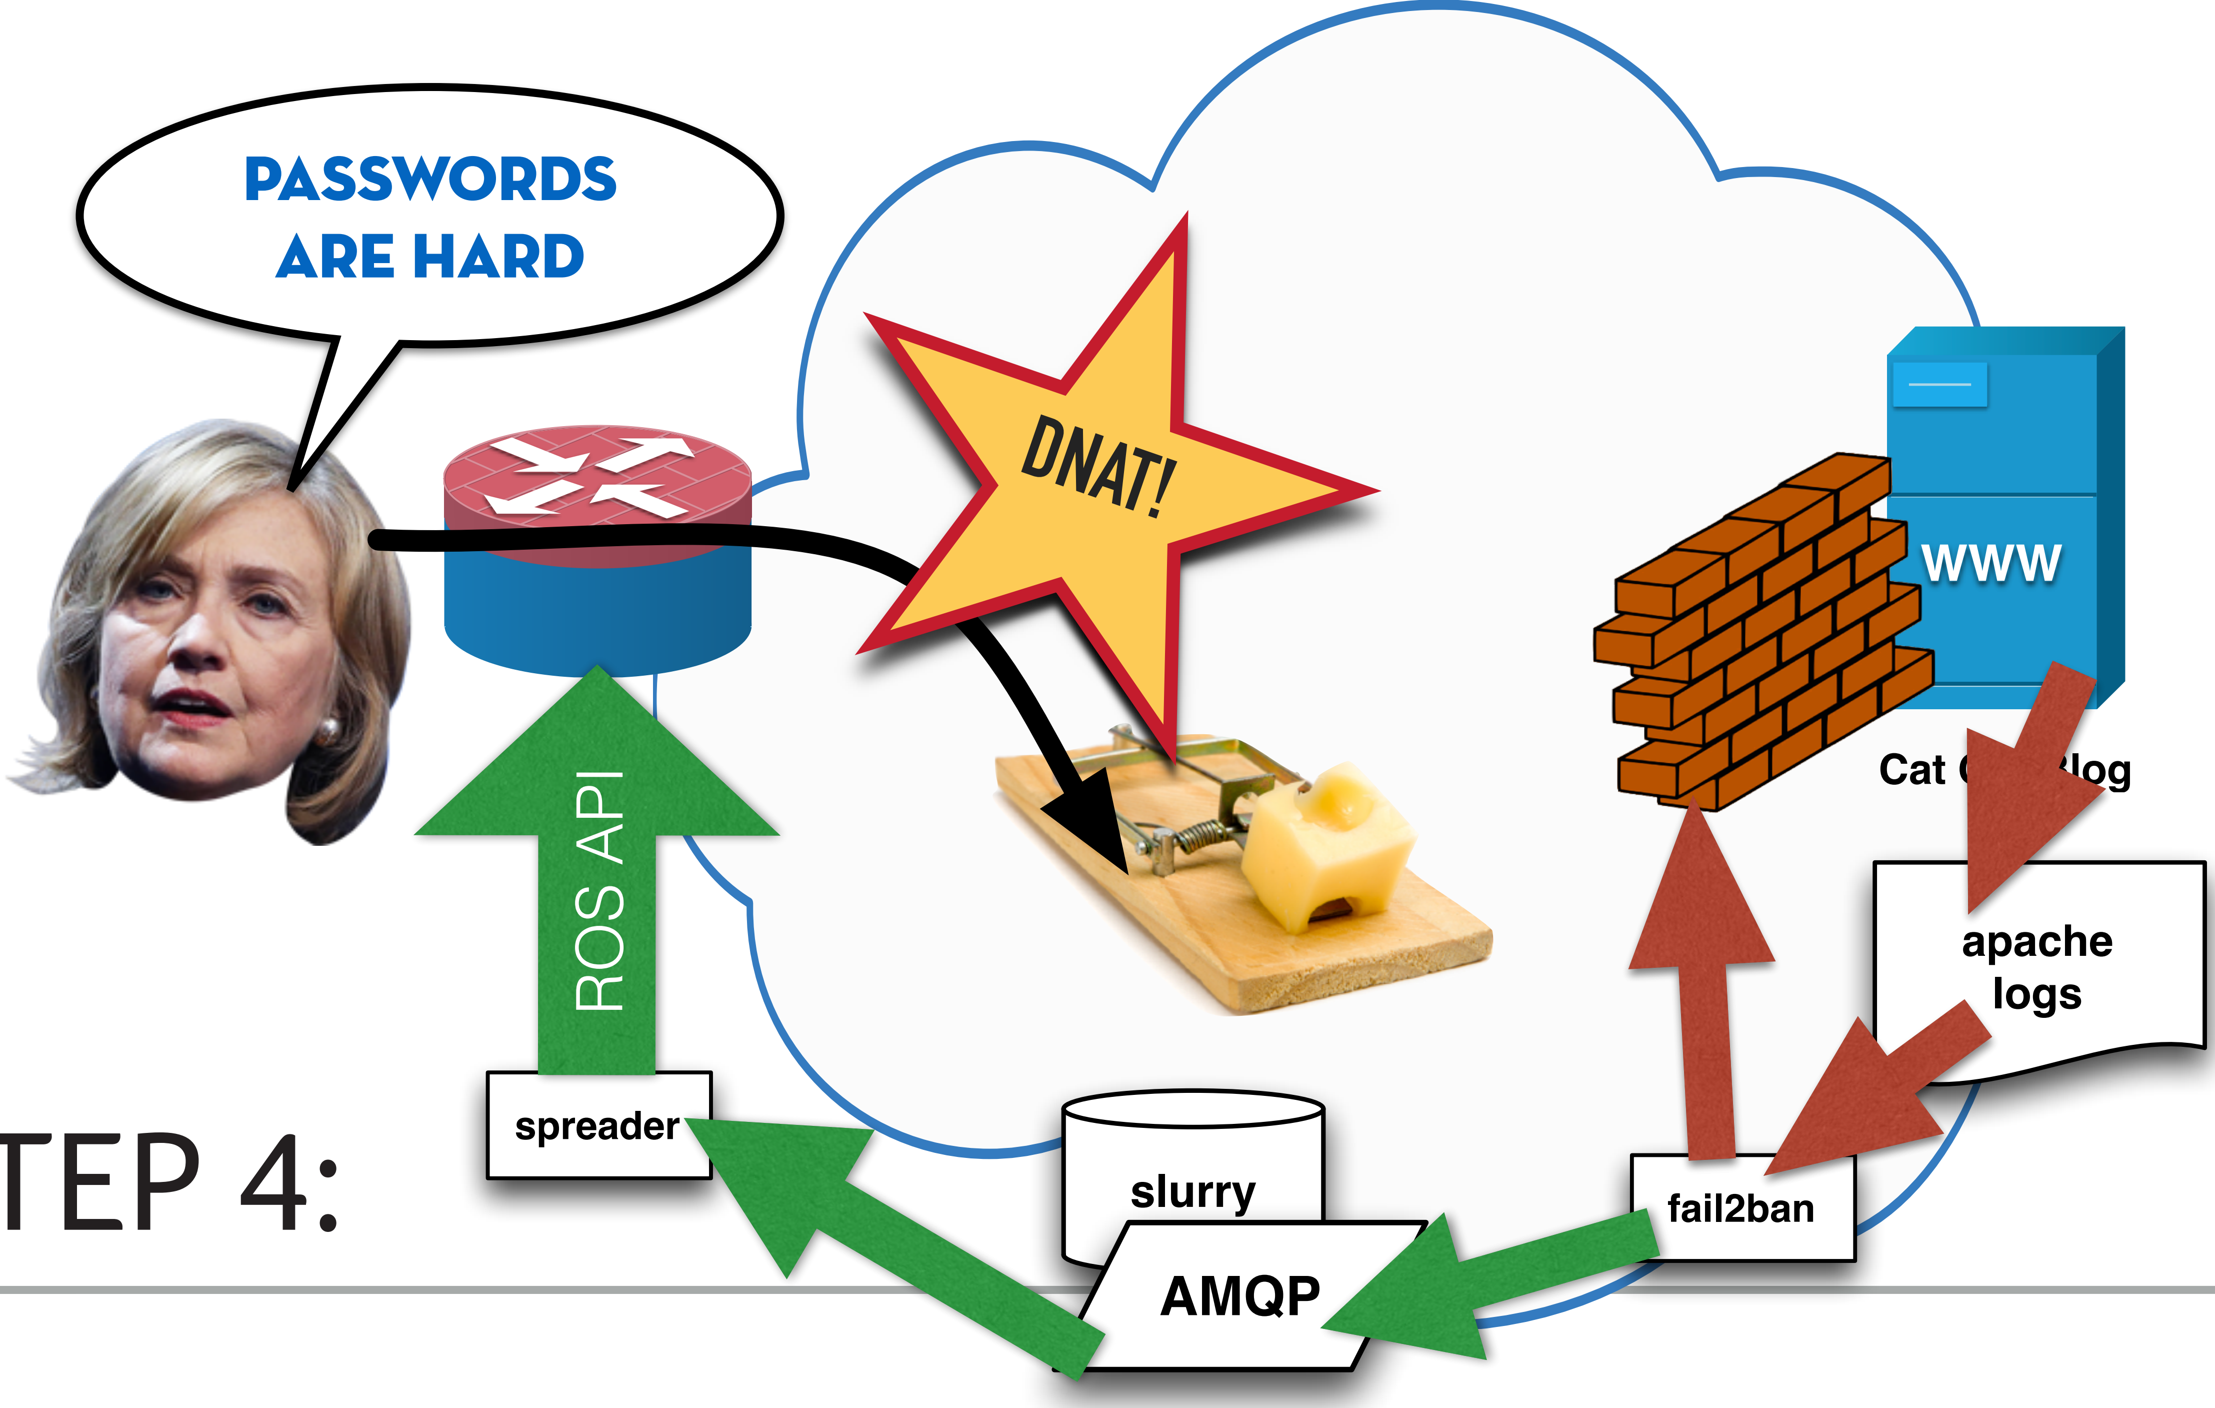

STEP 4:

FALSE POSITIVES

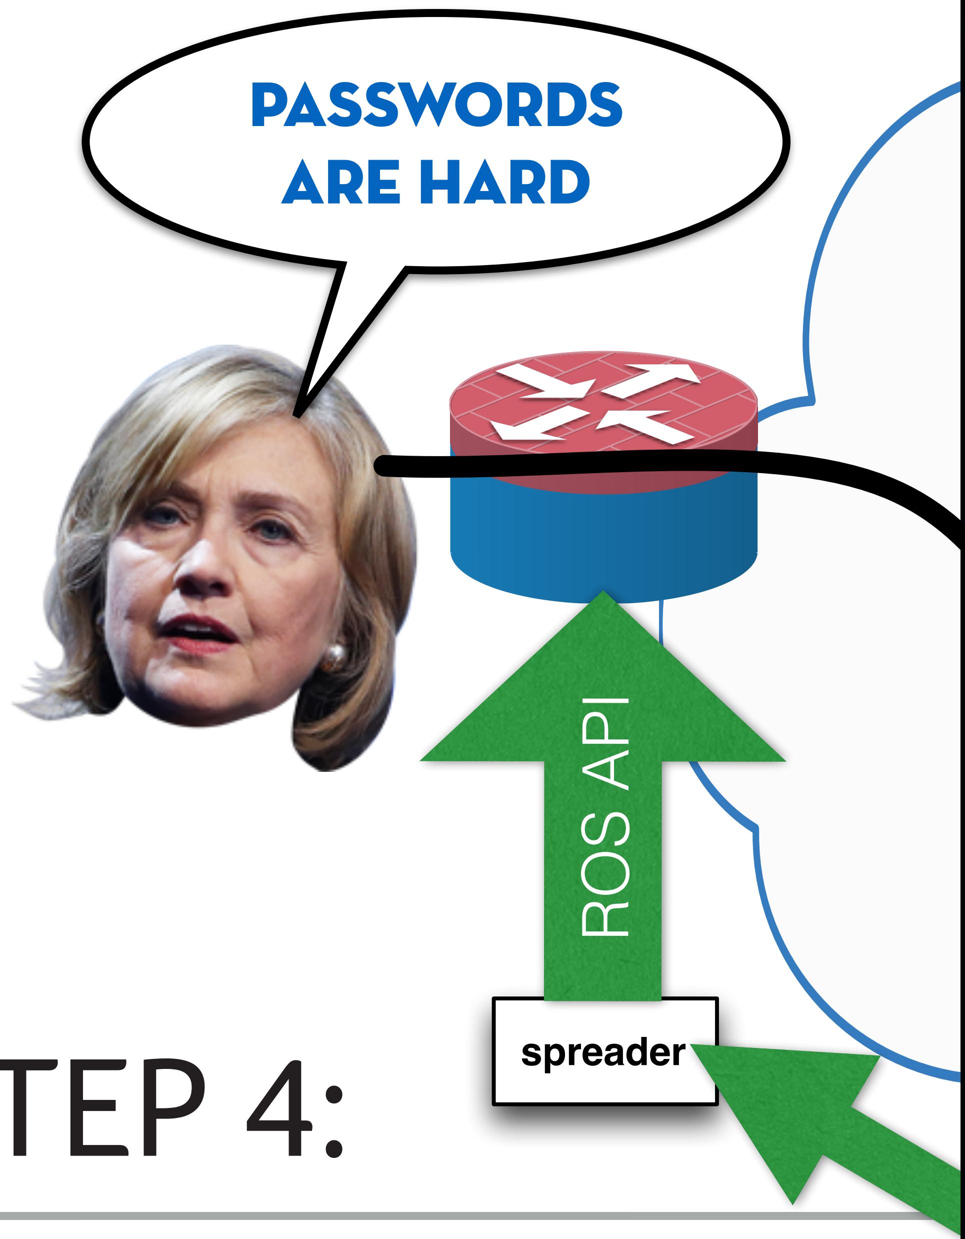

STEP 4:

FALSE P

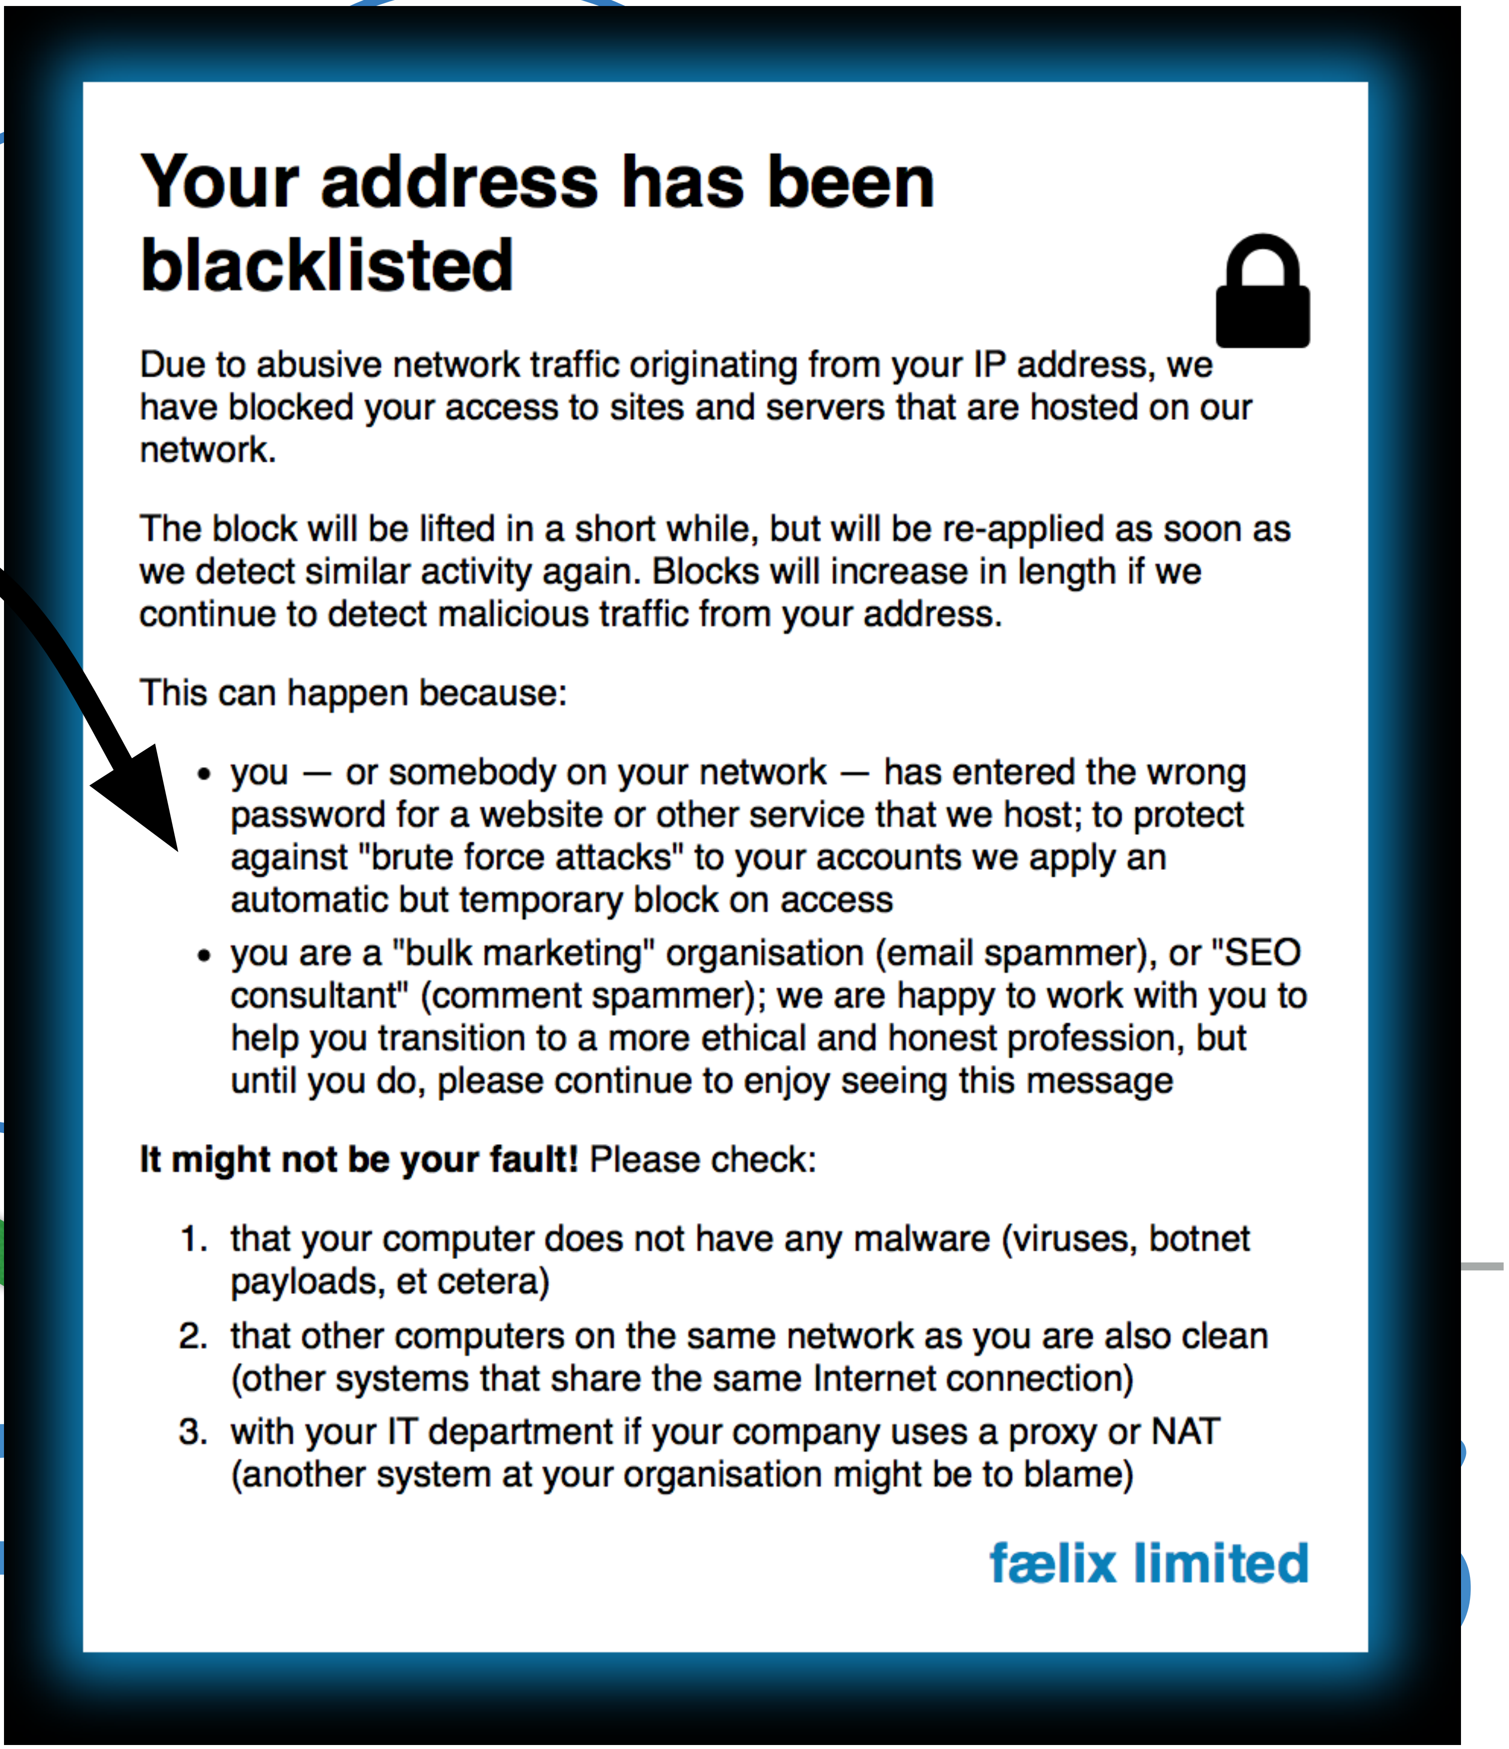

Your address has been blacklisted

Due to abusive network traffic originating from your IP address, we have blocked your access to sites and servers that are hosted on our network.

The block will be lifted in a short while, but will be re-applied as soon as we detect similar activity again. Blocks will increase in length if we continue to detect malicious traffic from your address.

This can happen because:

- you — or somebody on your network — has entered the wrong password for a website or other service that we host; to protect against "brute force attacks" to your accounts we apply an automatic but temporary block on access
- you are a "bulk marketing" organisation (email spammer), or "SEO consultant" (comment spammer); we are happy to work with you to help you transition to a more ethical and honest profession, but until you do, please continue to enjoy seeing this message

It might not be your fault! Please check:

1. that your computer does not have any malware (viruses, botnet payloads, et cetera)
2. that other computers on the same network as you are also clean (other systems that share the same Internet connection)
3. with your IT department if your company uses a proxy or NAT (another system at your organisation might be to blame)

fælix limited

DESTINATION NAT

- Send bad traffic to a VM serving the "blocked" message:
  - **/ip firewall nat src-address-list=shitpit action=dst-nat**

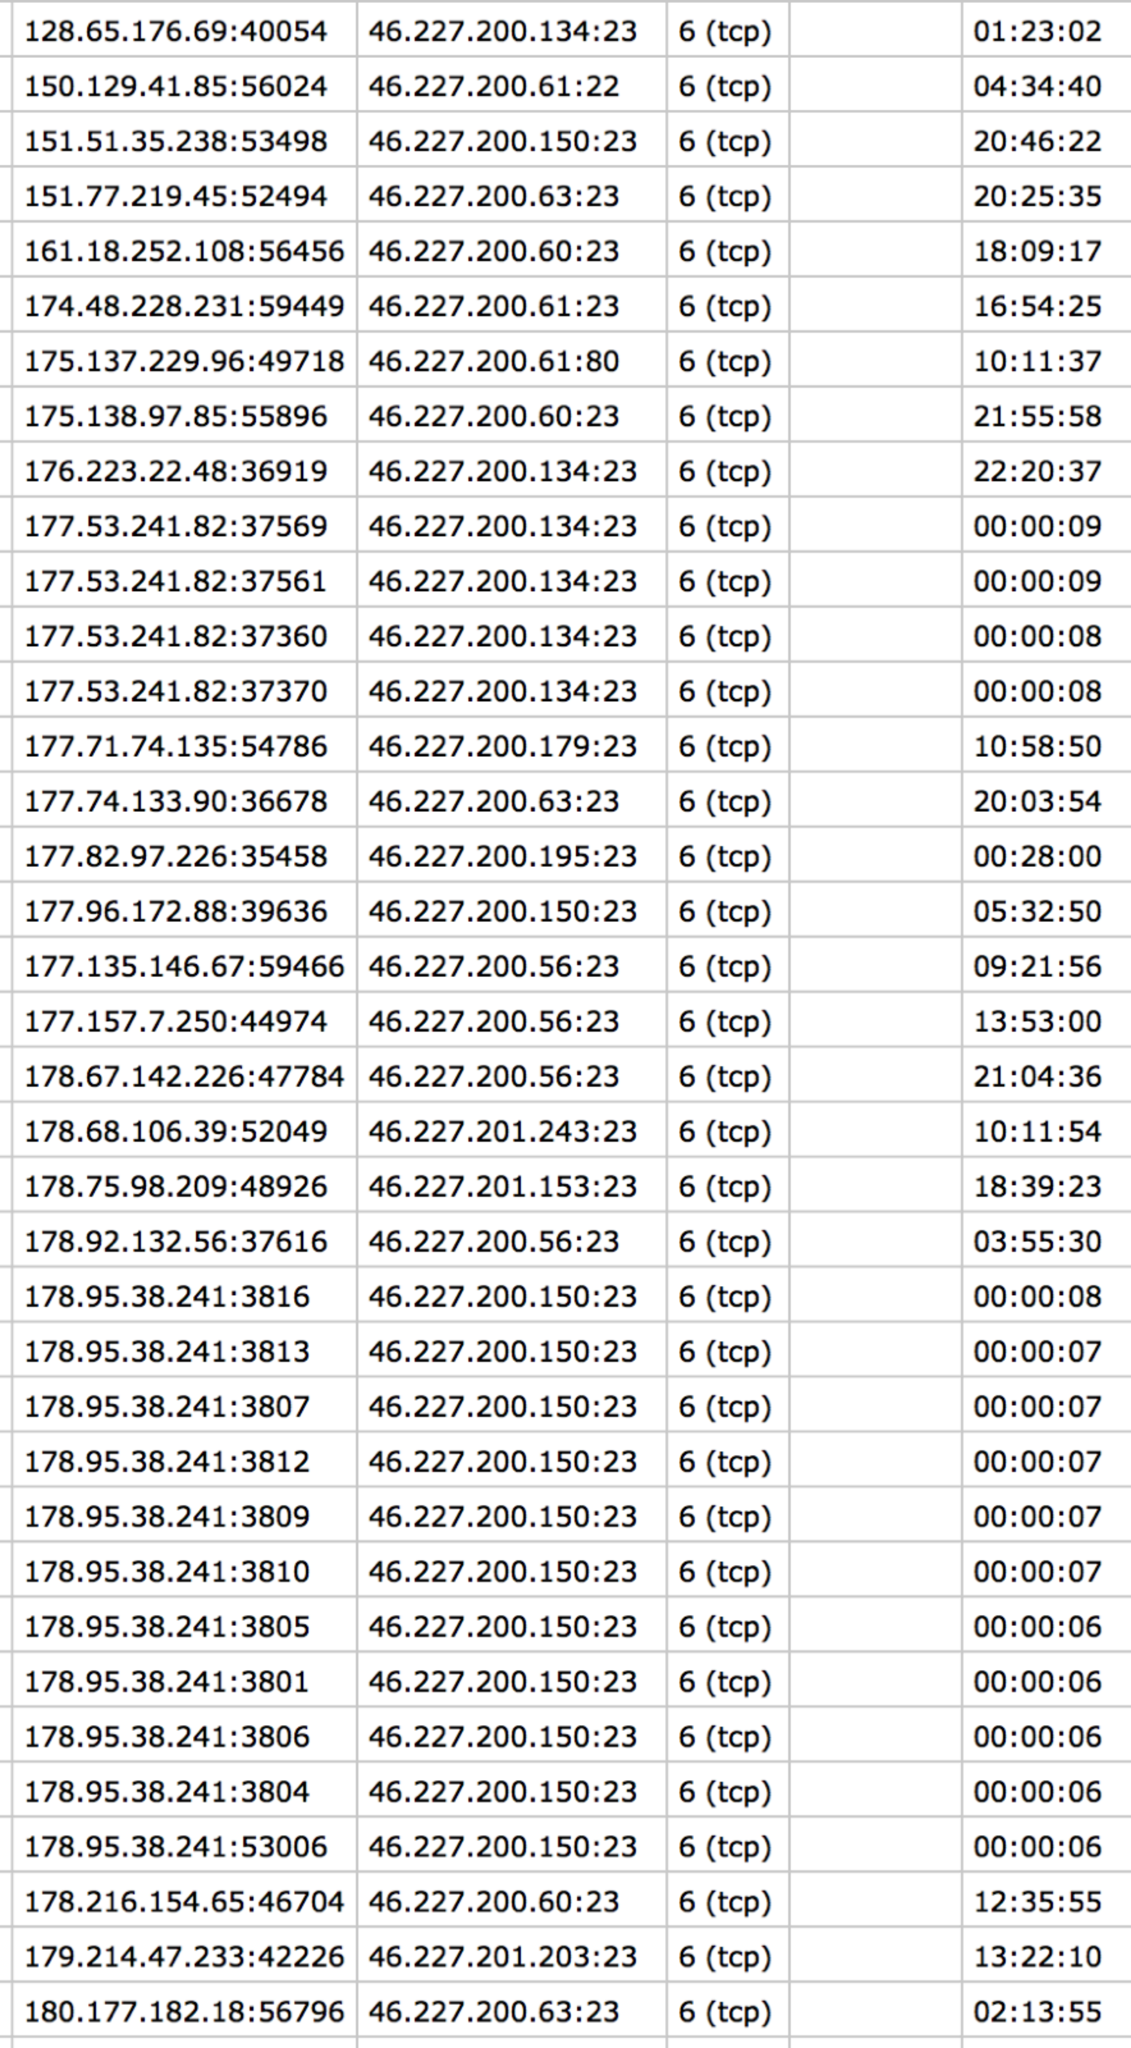

128.65.176.69:40054	46.227.200.134:23	6 (tcp)		01:23:02
150.129.41.85:56024	46.227.200.61:22	6 (tcp)		04:34:40
151.51.35.238:53498	46.227.200.150:23	6 (tcp)		20:46:22
151.77.219.45:52494	46.227.200.63:23	6 (tcp)		20:25:35
161.18.252.108:56456	46.227.200.60:23	6 (tcp)		18:09:17
174.48.228.231:59449	46.227.200.61:23	6 (tcp)		16:54:25
175.137.229.96:49718	46.227.200.61:80	6 (tcp)		10:11:37
175.138.97.85:55896	46.227.200.60:23	6 (tcp)		21:55:58
176.223.22.48:36919	46.227.200.134:23	6 (tcp)		22:20:37
177.53.241.82:37569	46.227.200.134:23	6 (tcp)		00:00:09
177.53.241.82:37561	46.227.200.134:23	6 (tcp)		00:00:09
177.53.241.82:37360	46.227.200.134:23	6 (tcp)		00:00:08
177.53.241.82:37370	46.227.200.134:23	6 (tcp)		00:00:08
177.71.74.135:54786	46.227.200.179:23	6 (tcp)		10:58:50
177.74.133.90:36678	46.227.200.63:23	6 (tcp)		20:03:54
177.82.97.226:35458	46.227.200.195:23	6 (tcp)		00:28:00
177.96.172.88:39636	46.227.200.150:23	6 (tcp)		05:32:50
177.135.146.67:59466	46.227.200.56:23	6 (tcp)		09:21:56
177.157.7.250:44974	46.227.200.56:23	6 (tcp)		13:53:00
178.67.142.226:47784	46.227.200.56:23	6 (tcp)		21:04:36
178.68.106.39:52049	46.227.201.243:23	6 (tcp)		10:11:54
178.75.98.209:48926	46.227.201.153:23	6 (tcp)		18:39:23
178.92.132.56:37616	46.227.200.56:23	6 (tcp)		03:55:30
178.95.38.241:3816	46.227.200.150:23	6 (tcp)		00:00:08
178.95.38.241:3813	46.227.200.150:23	6 (tcp)		00:00:07
178.95.38.241:3807	46.227.200.150:23	6 (tcp)		00:00:07
178.95.38.241:3812	46.227.200.150:23	6 (tcp)		00:00:07
178.95.38.241:3809	46.227.200.150:23	6 (tcp)		00:00:07
178.95.38.241:3810	46.227.200.150:23	6 (tcp)		00:00:07
178.95.38.241:3805	46.227.200.150:23	6 (tcp)		00:00:06
178.95.38.241:3801	46.227.200.150:23	6 (tcp)		00:00:06
178.95.38.241:3806	46.227.200.150:23	6 (tcp)		00:00:06
178.95.38.241:3804	46.227.200.150:23	6 (tcp)		00:00:06
178.95.38.241:53006	46.227.200.150:23	6 (tcp)		00:00:06
178.216.154.65:46704	46.227.200.60:23	6 (tcp)		12:35:55
179.214.47.233:42226	46.227.201.203:23	6 (tcp)		13:22:10
180.177.182.18:56796	46.227.200.63:23	6 (tcp)		02:13:55

CONN TRACK!

tl;dr: ah, crap

BLOCKING AT THE PROVIDER EDGE

- Lots of flows, lots of PPS, lots of attacking addresses
  - **/ip firewall filter** uses each set of rules sequentially = O(n)
  - **/ip firewall address-list** is a hash-table ≈ O(1)
- Using AMQP to get addresses added to block lists on all routers in three data-centres
  - We already had RabbitMQ across our network for other infrastructure needs

BLOCKING AT THE PROVIDER EDGE

- Lots of flows…
  - …so use a **mangle** rule so routers only track bad traffic?
  - No! We want to build something we can understand.

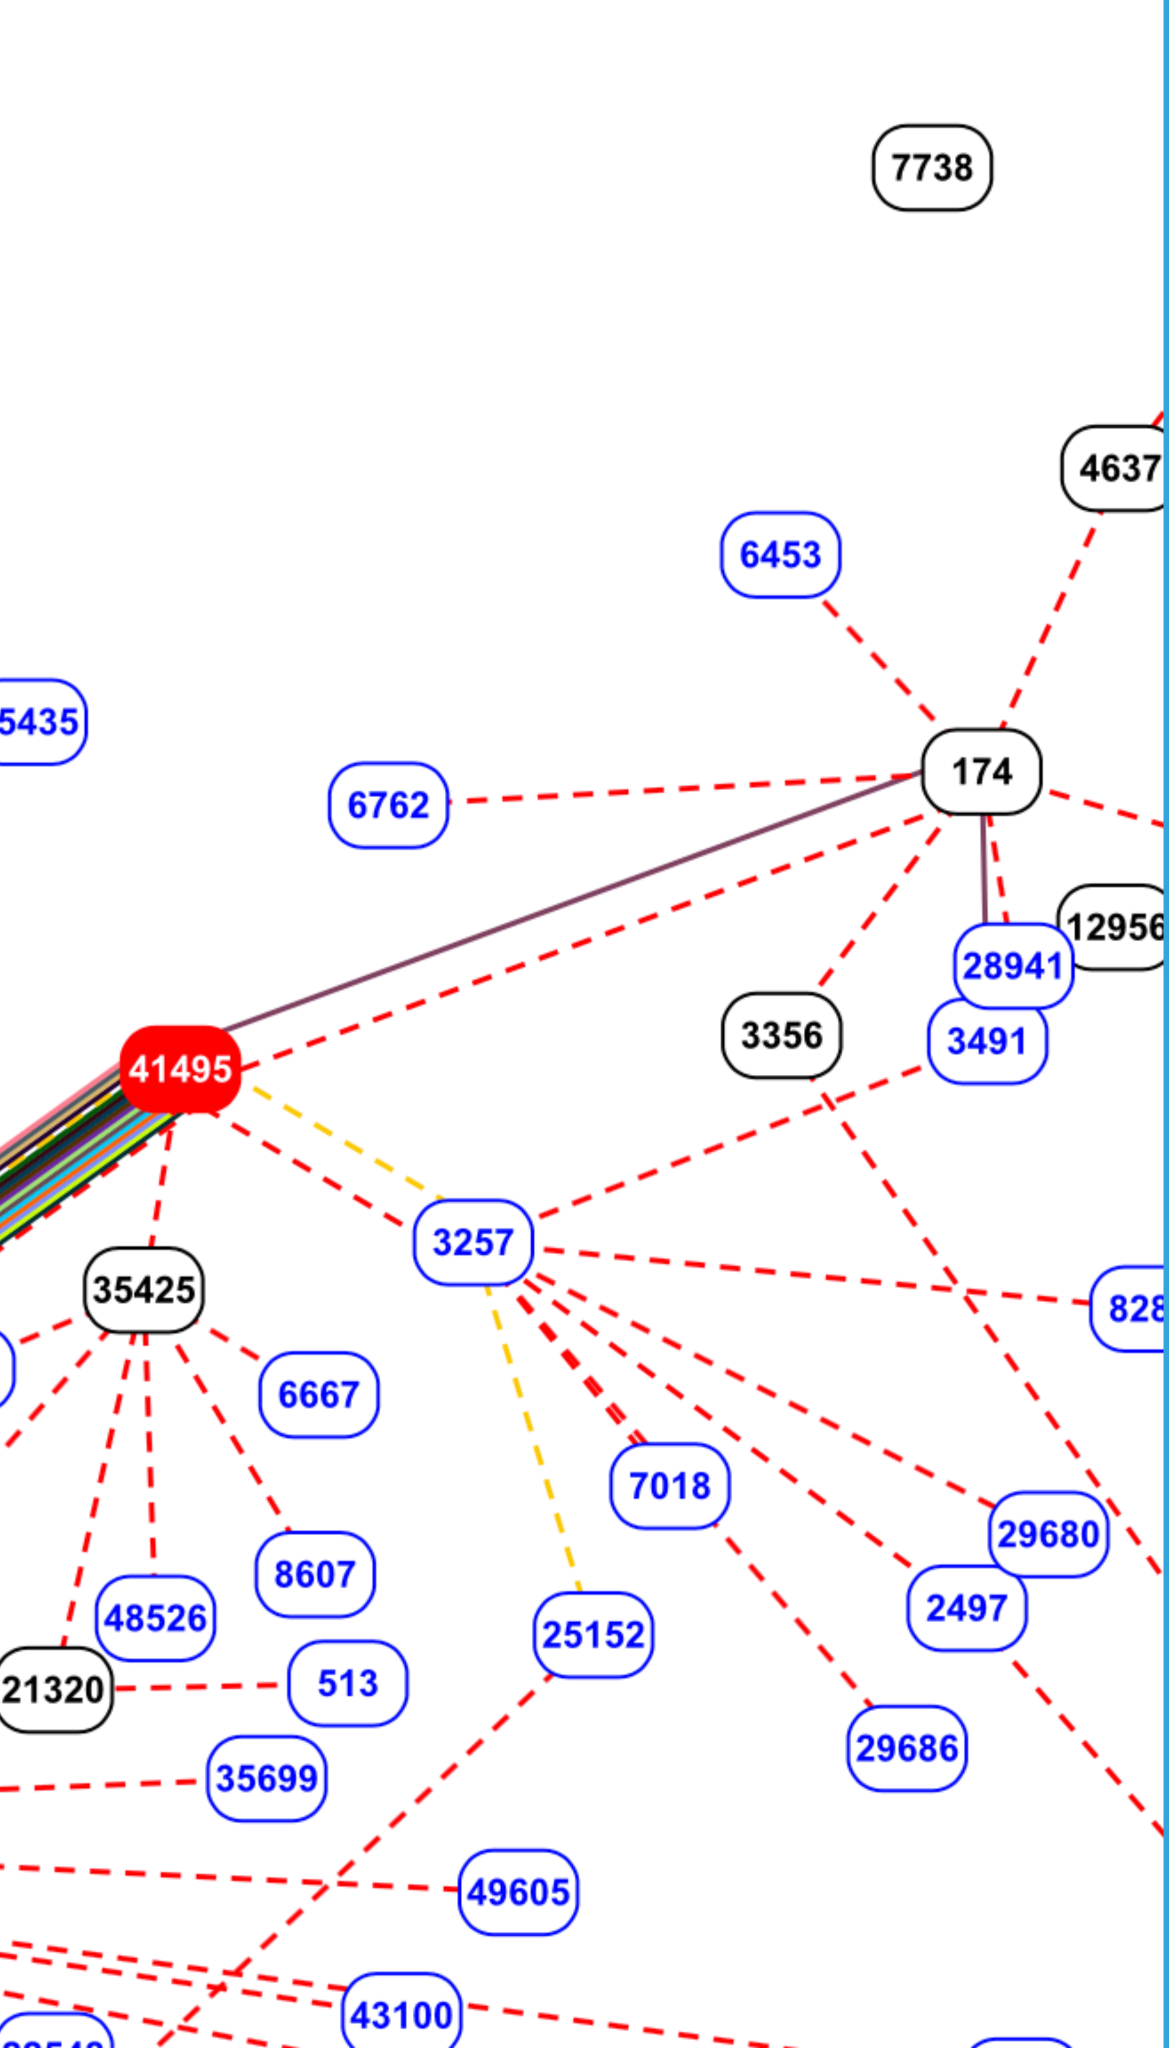

MULTI-HOMED!

tl;dr: ah, crap^2

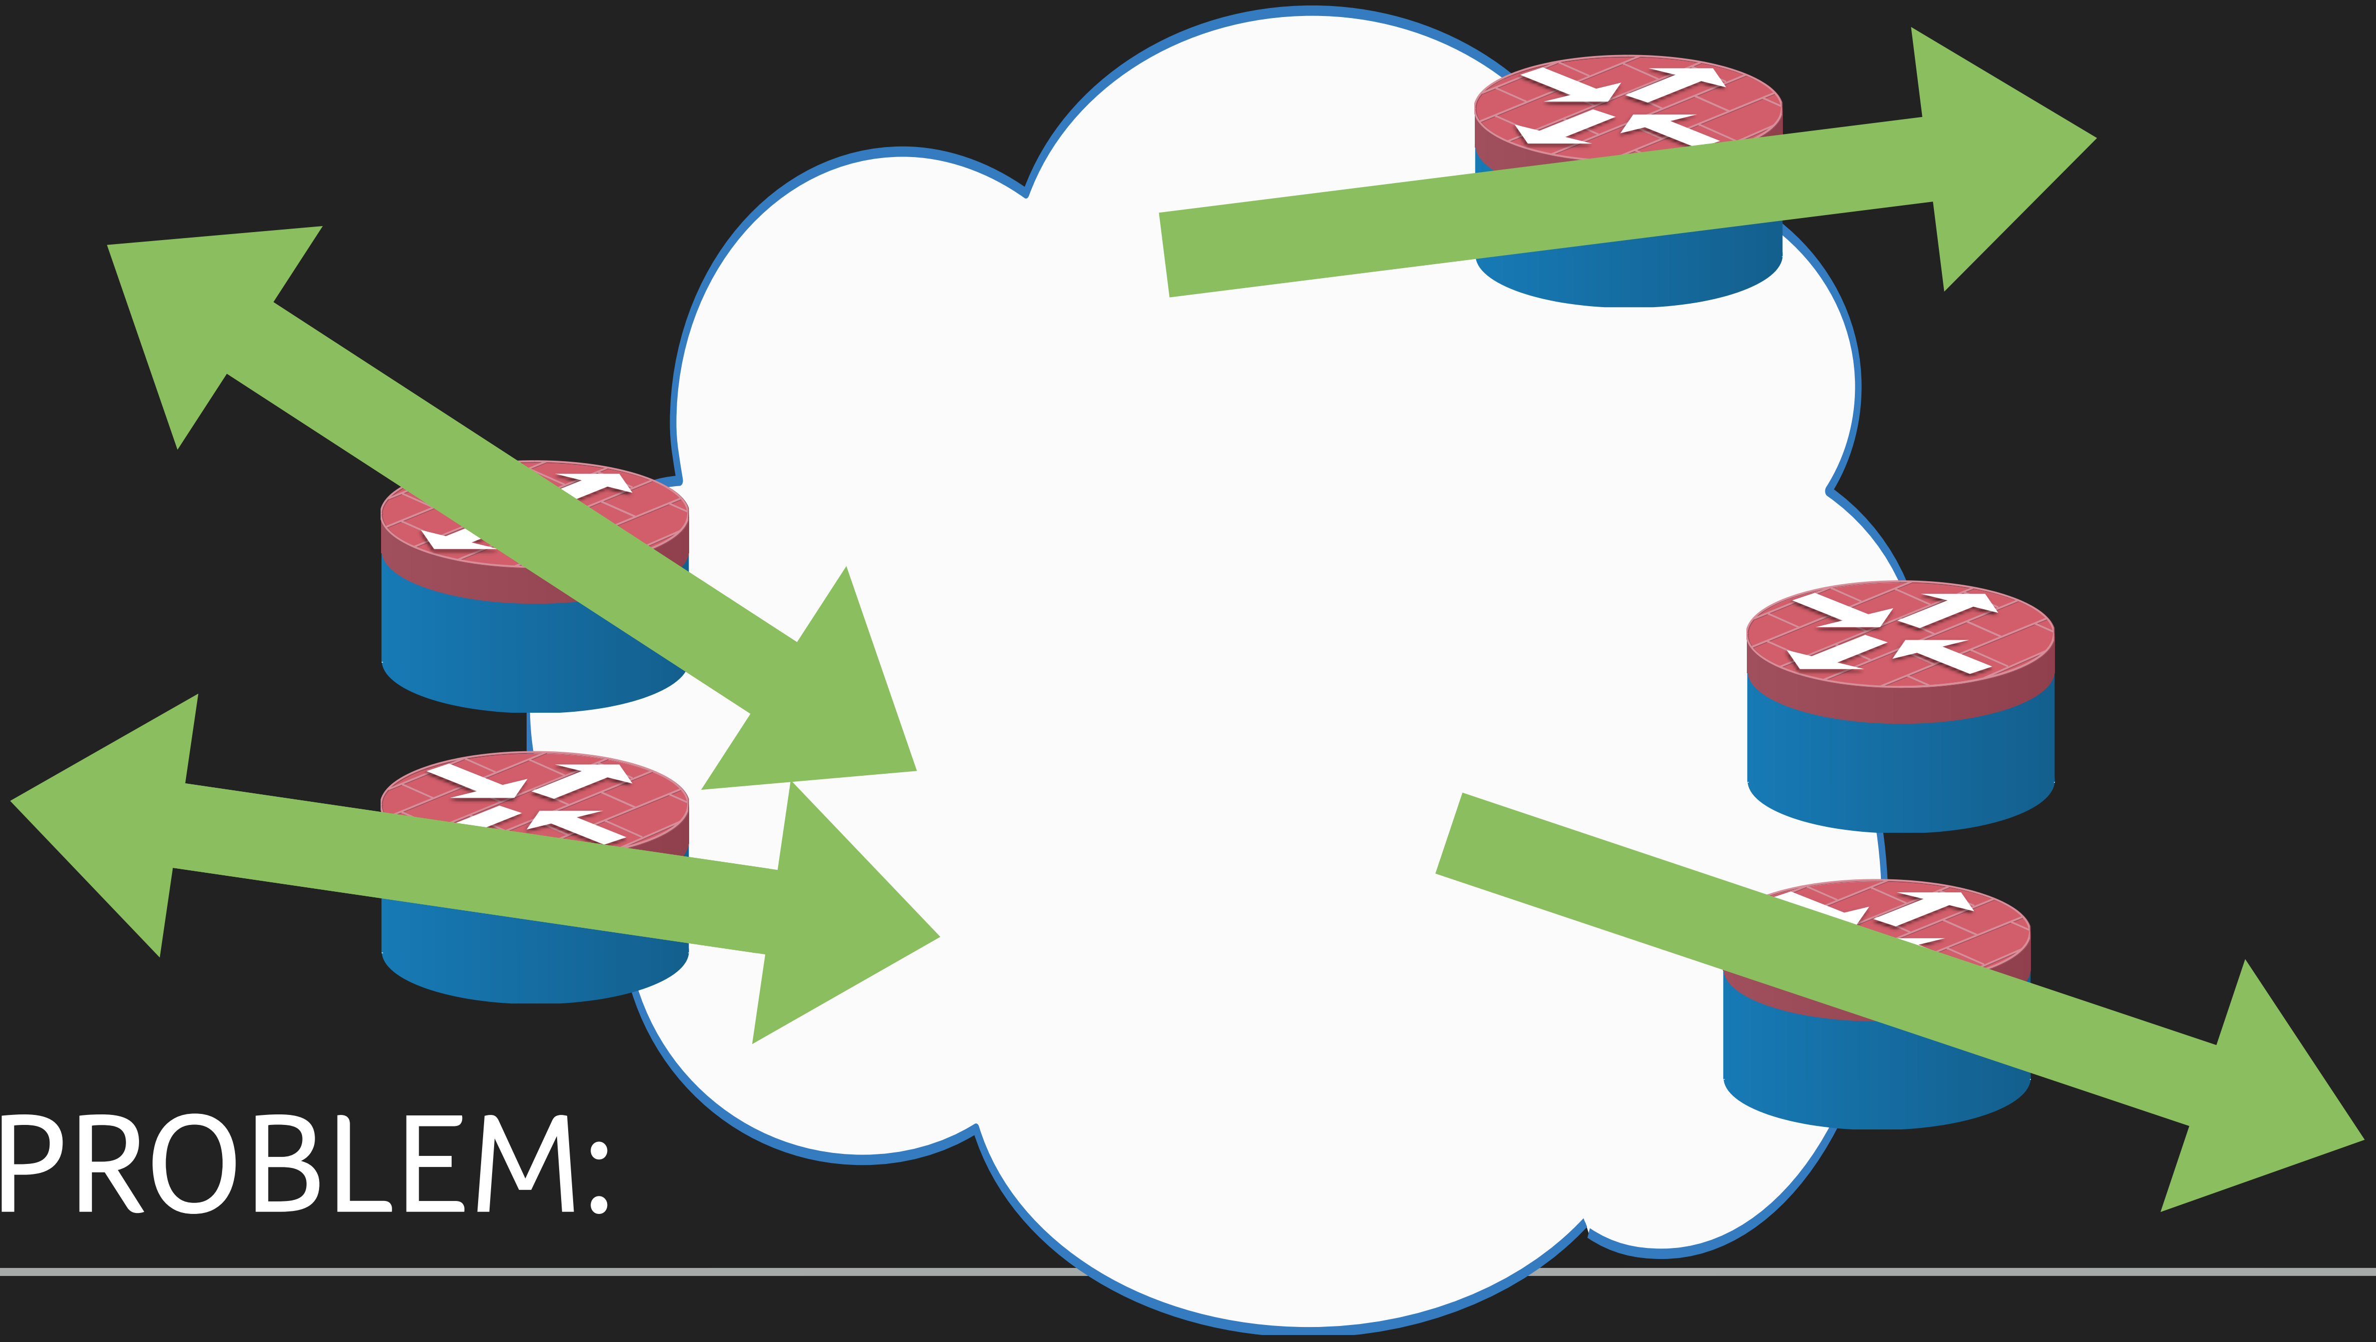

PROBLEM:

TRAFFIC HAS MULTIPLE PATHS IN AND OUT OF OUR NETWORK

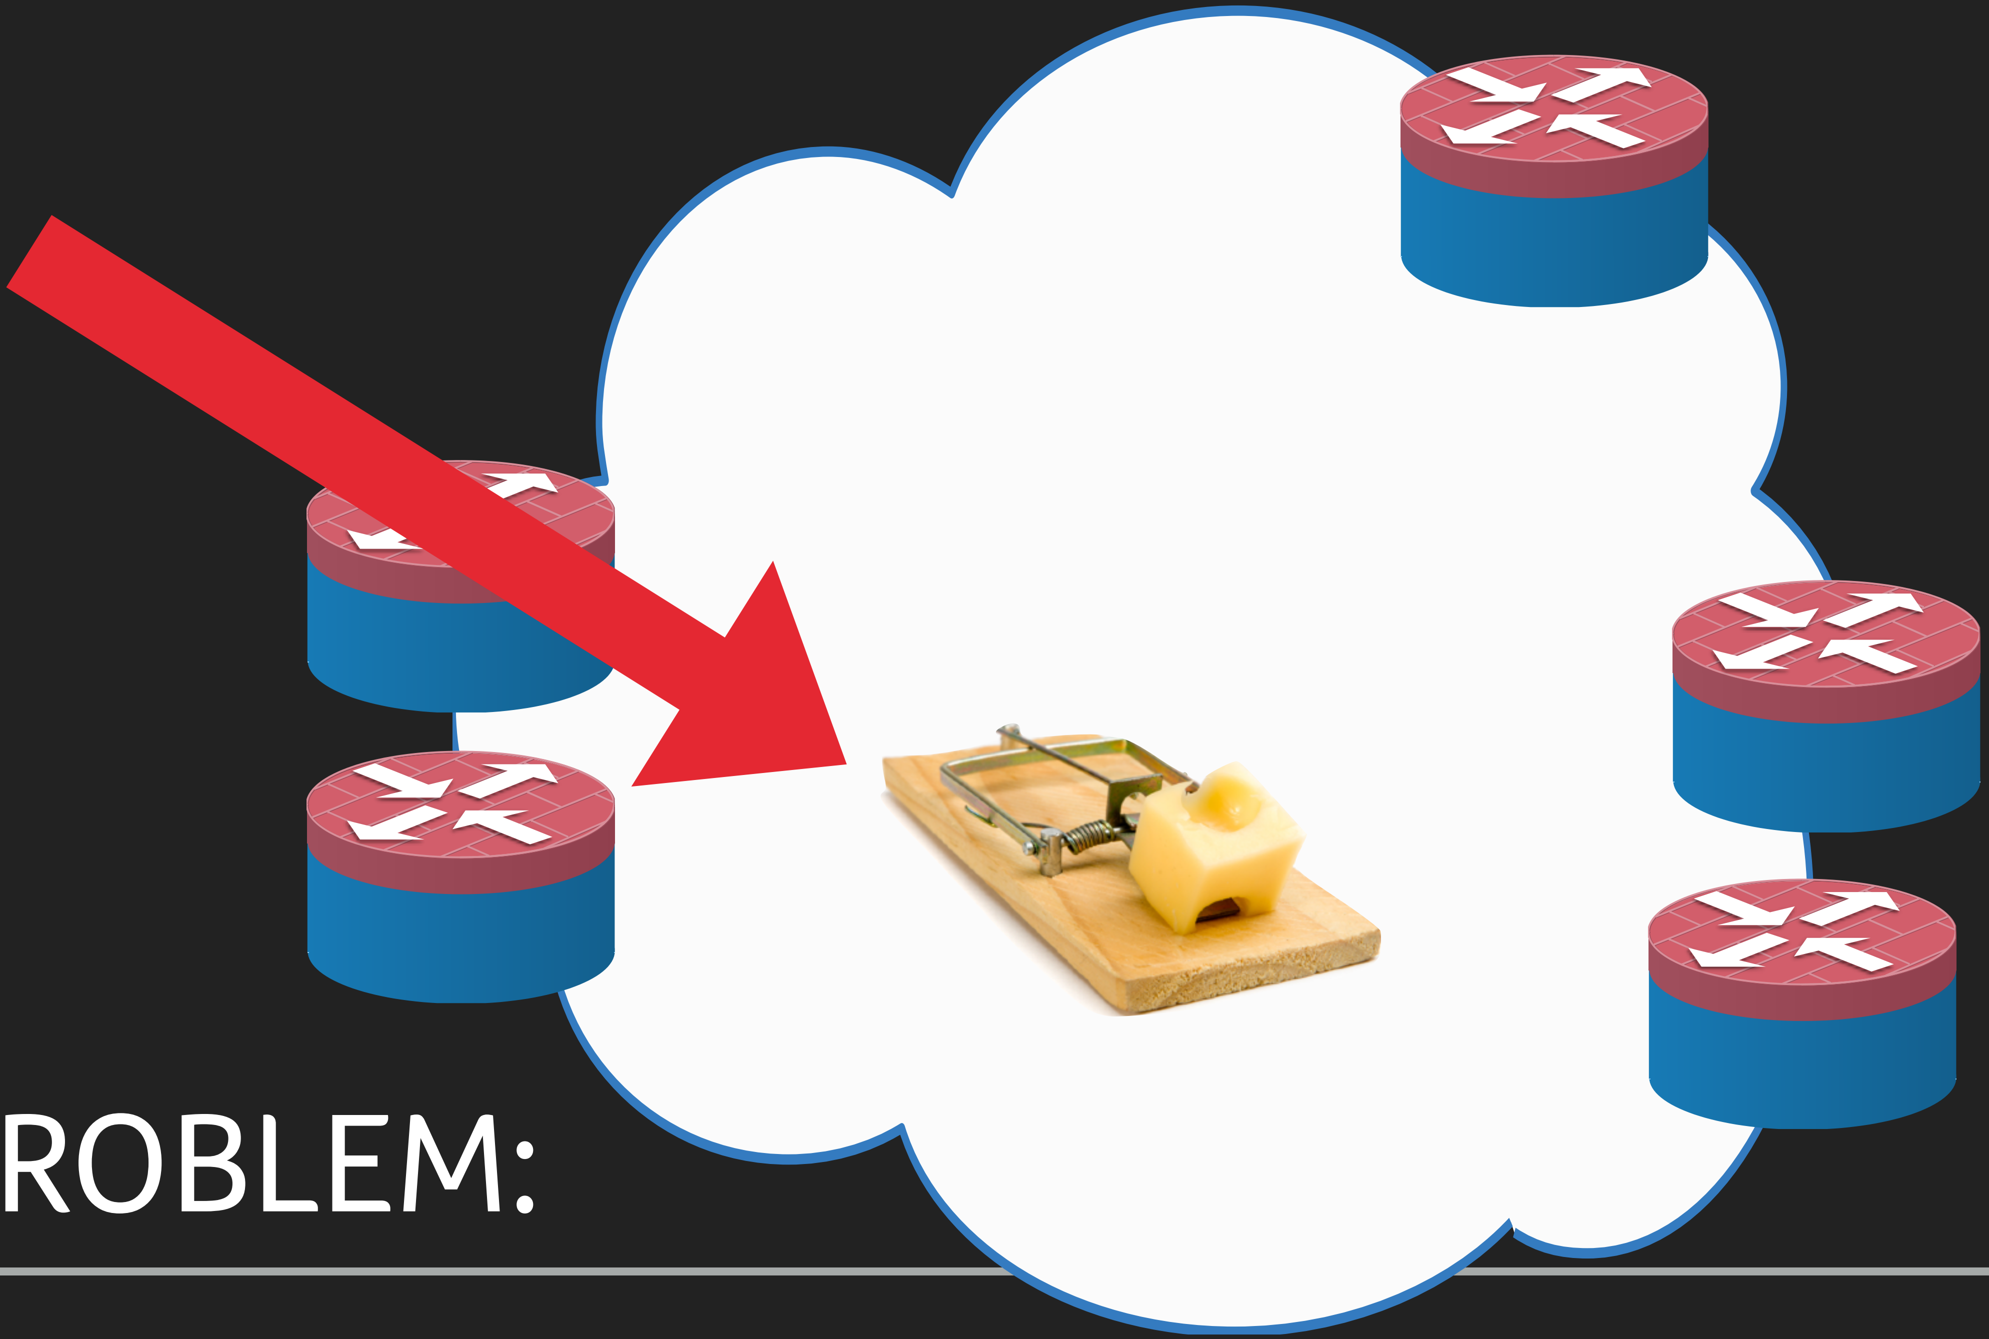

PROBLEM:

TRAFFIC HAS MULTIPLE PATHS IN AND OUT OF OUR NETWORK

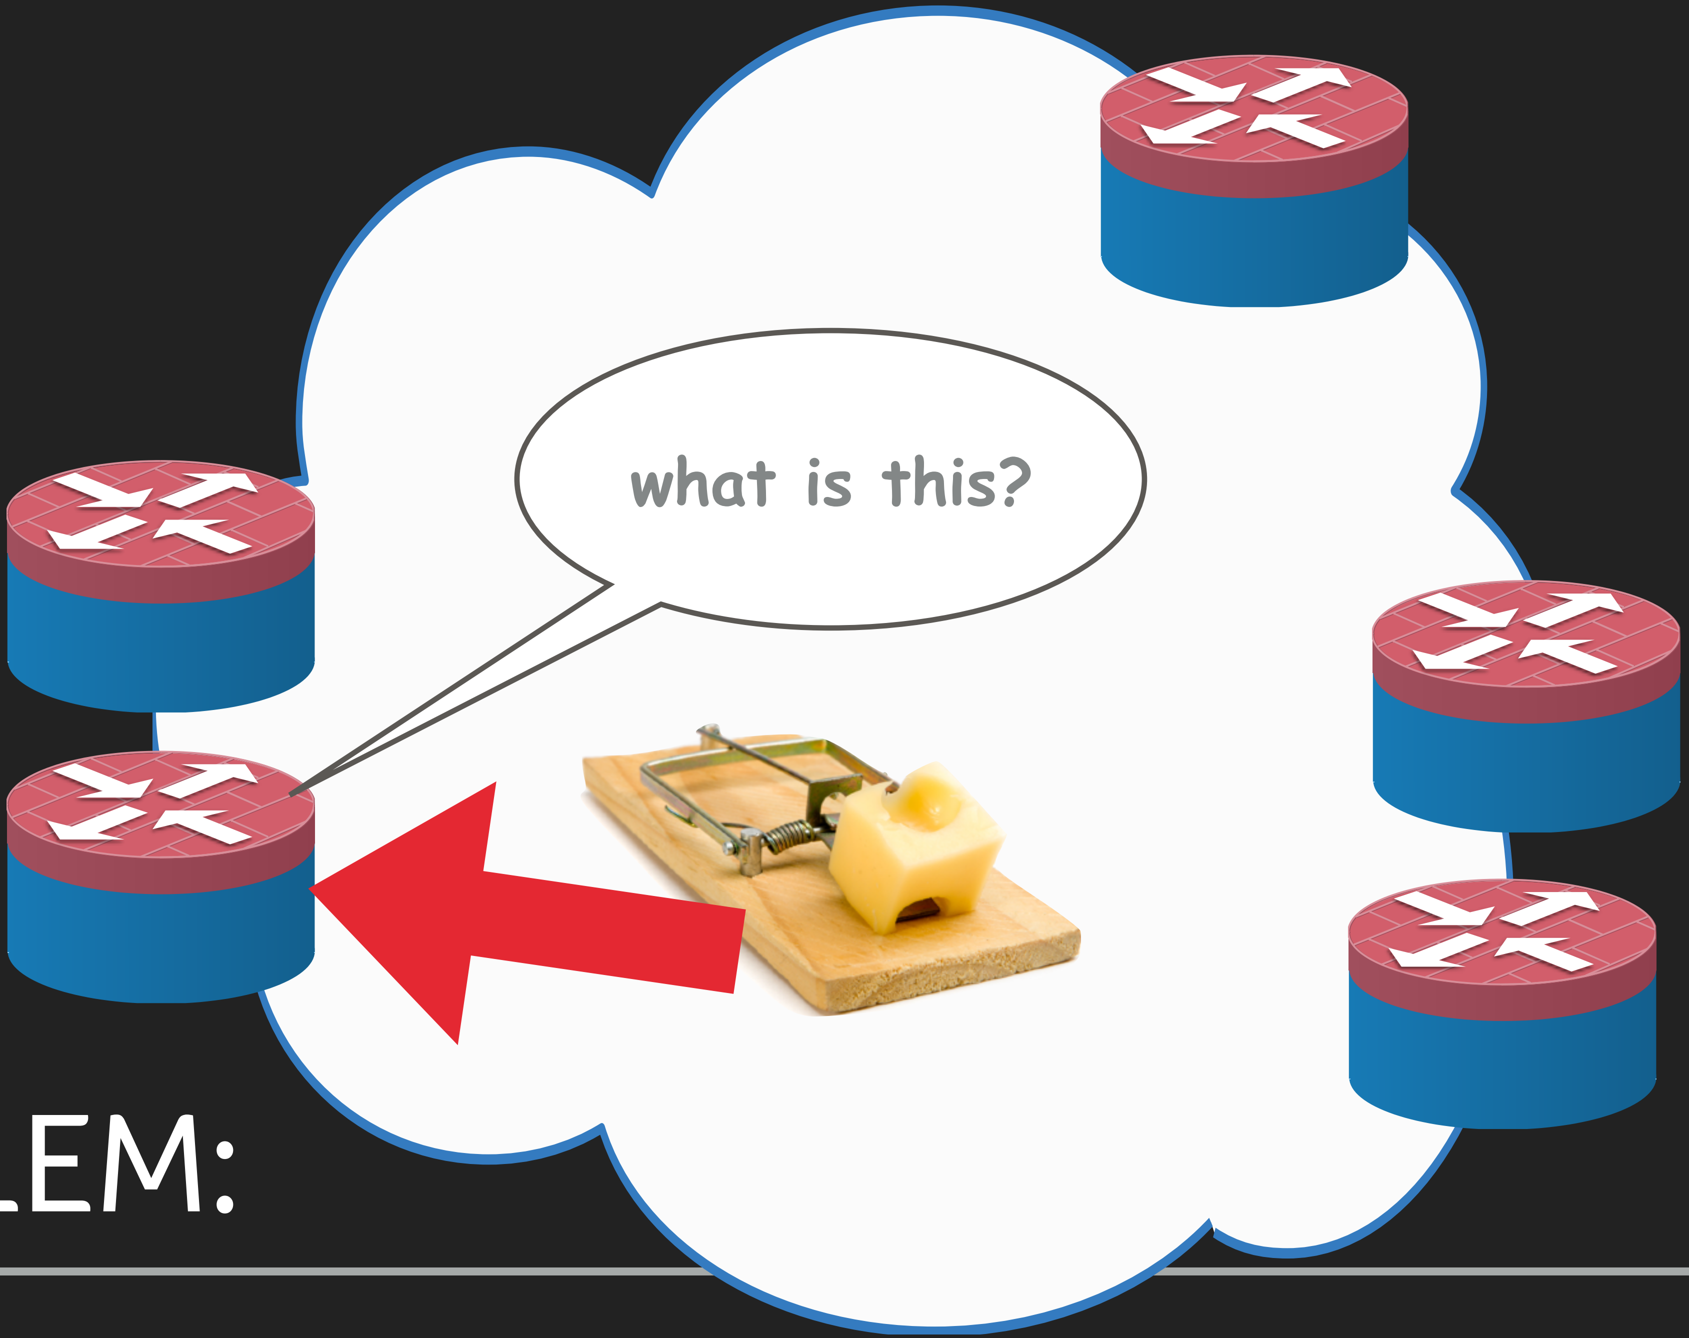

PROBLEM:

TRAFFIC HAS MULTIPLE PATHS IN AND OUT OF OUR NETWORK

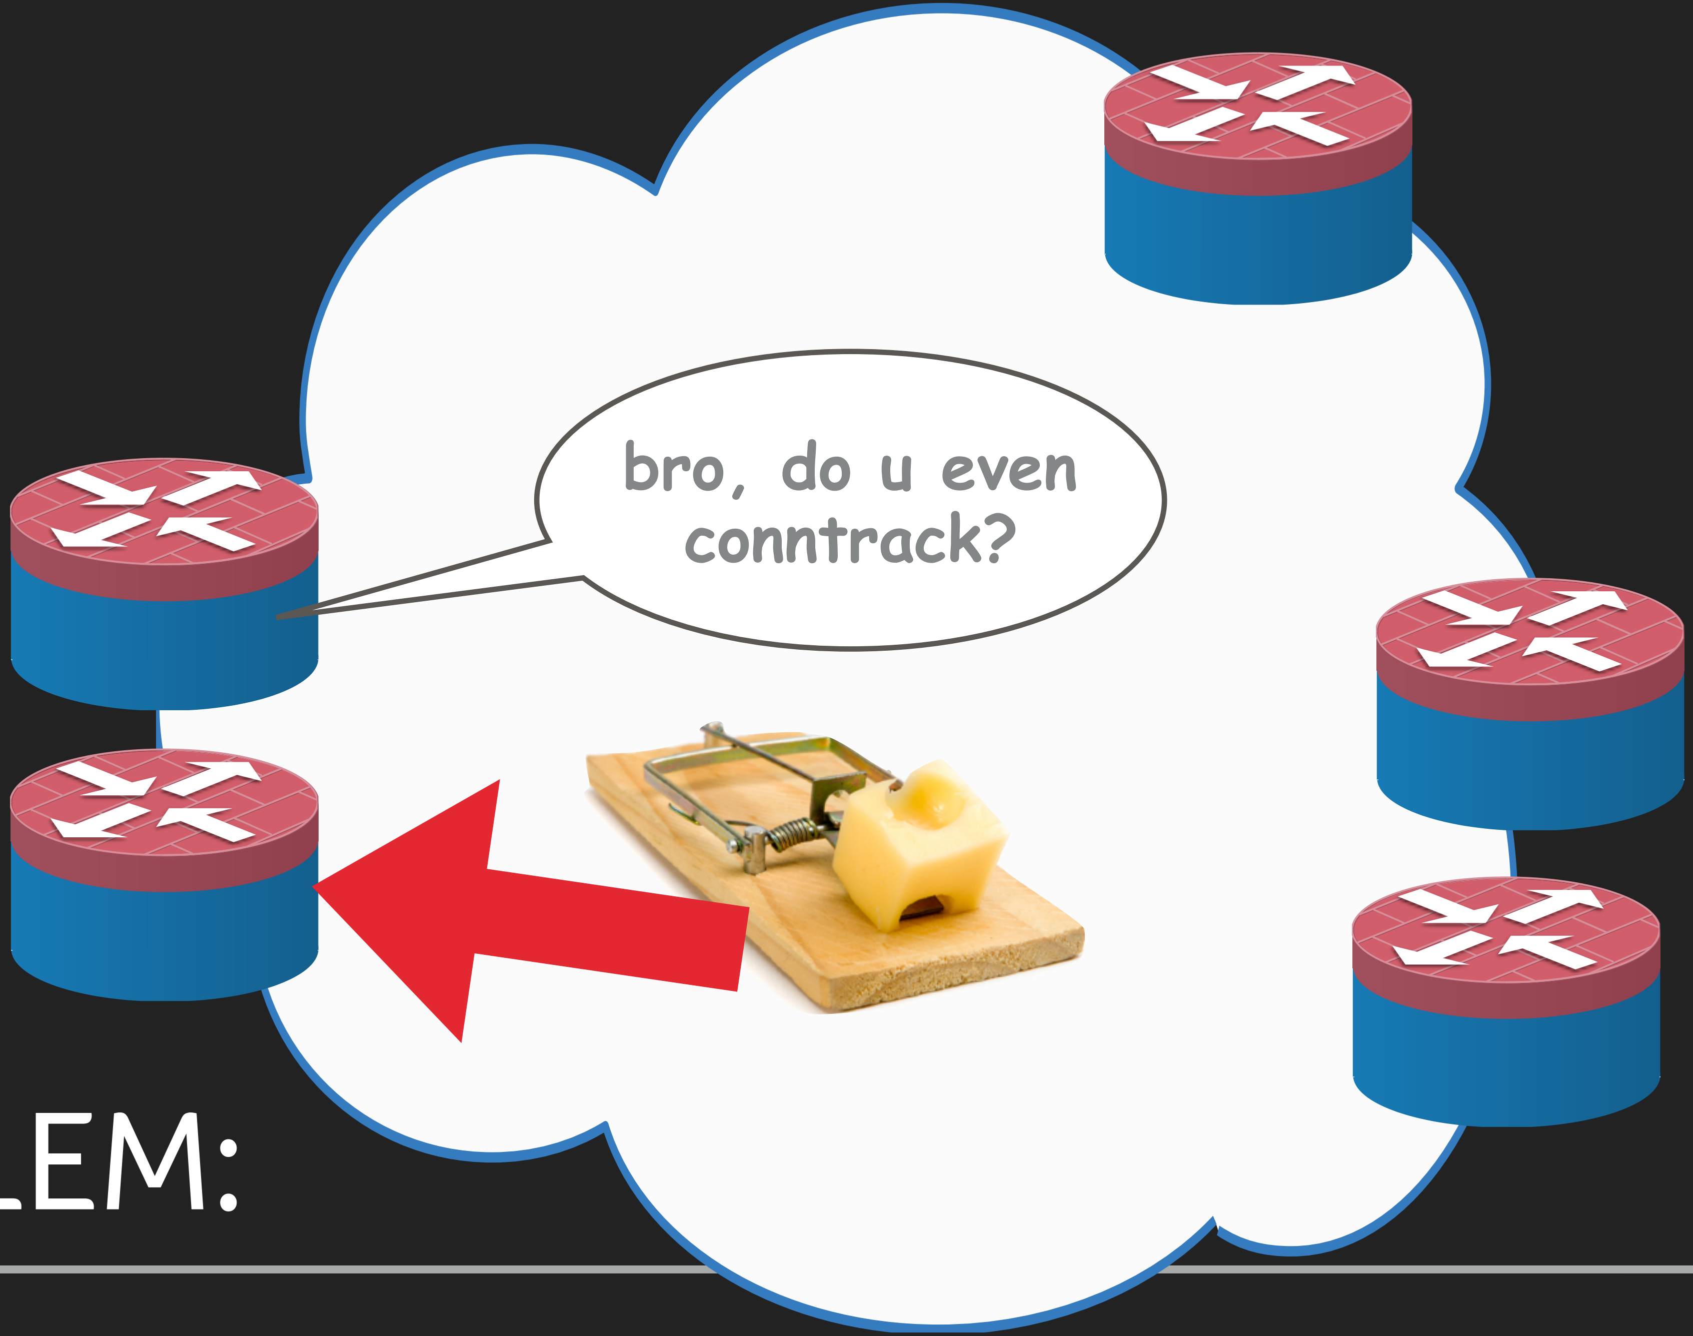

PROBLEM:

TRAFFIC HAS MULTIPLE PATHS IN AND OUT OF OUR NETWORK

WON'T CONNTRACK, CAN'T NAT

- Lots of flows
- Can't share conntrack across RouterOS devices
  - Would be nice for VRRP-type HA default gateways?
  - We don't want to even if we could: lots of flows!
    - And don't want to mangle to ignore good flows…
    - …and mangle to make return traffic go the right way.
- **"Are we there yet!?"**

MULTIPLE ROUTING TABLES

- /ip route add gateway=203.0.113.113 routing-mark=shitpit
- /ip route rule add routing-mark=shitpit table=shitpit
- /ip firewall mangle add chain=prerouting passthrough=yes action=mark-routing new-routing-mark=shitpit src-address-list=shitpit
- /ip firewall address-list add list=shitpit address=192.0.2.69/32 timeout=1m

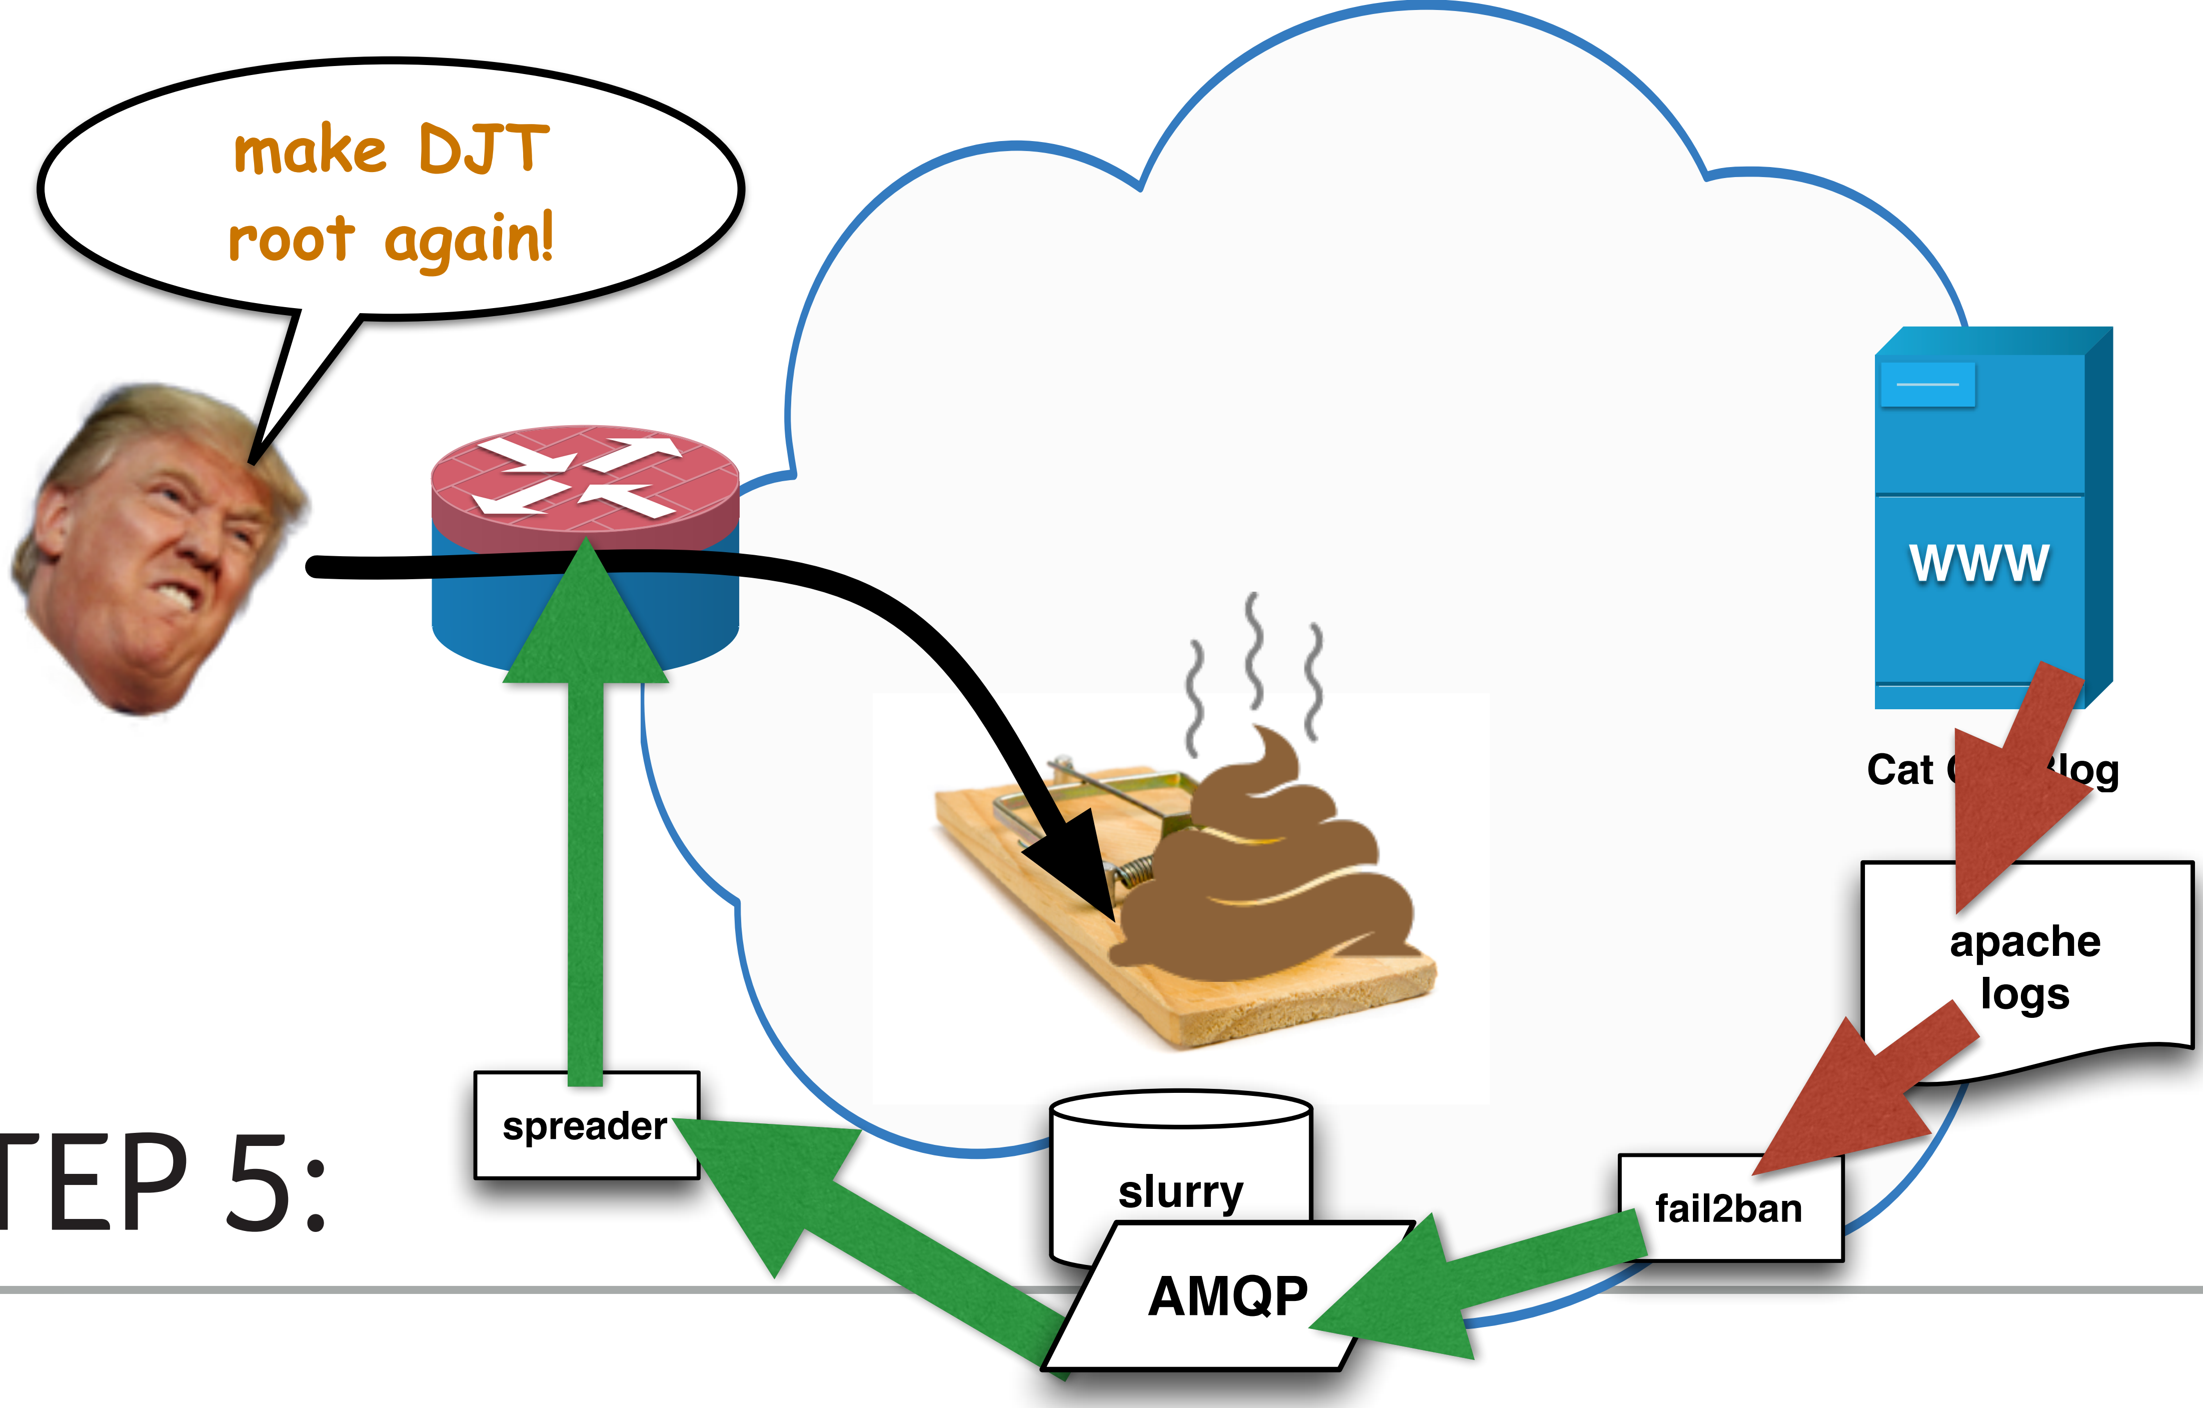

STEP 5:

/IP ROUTE RULE

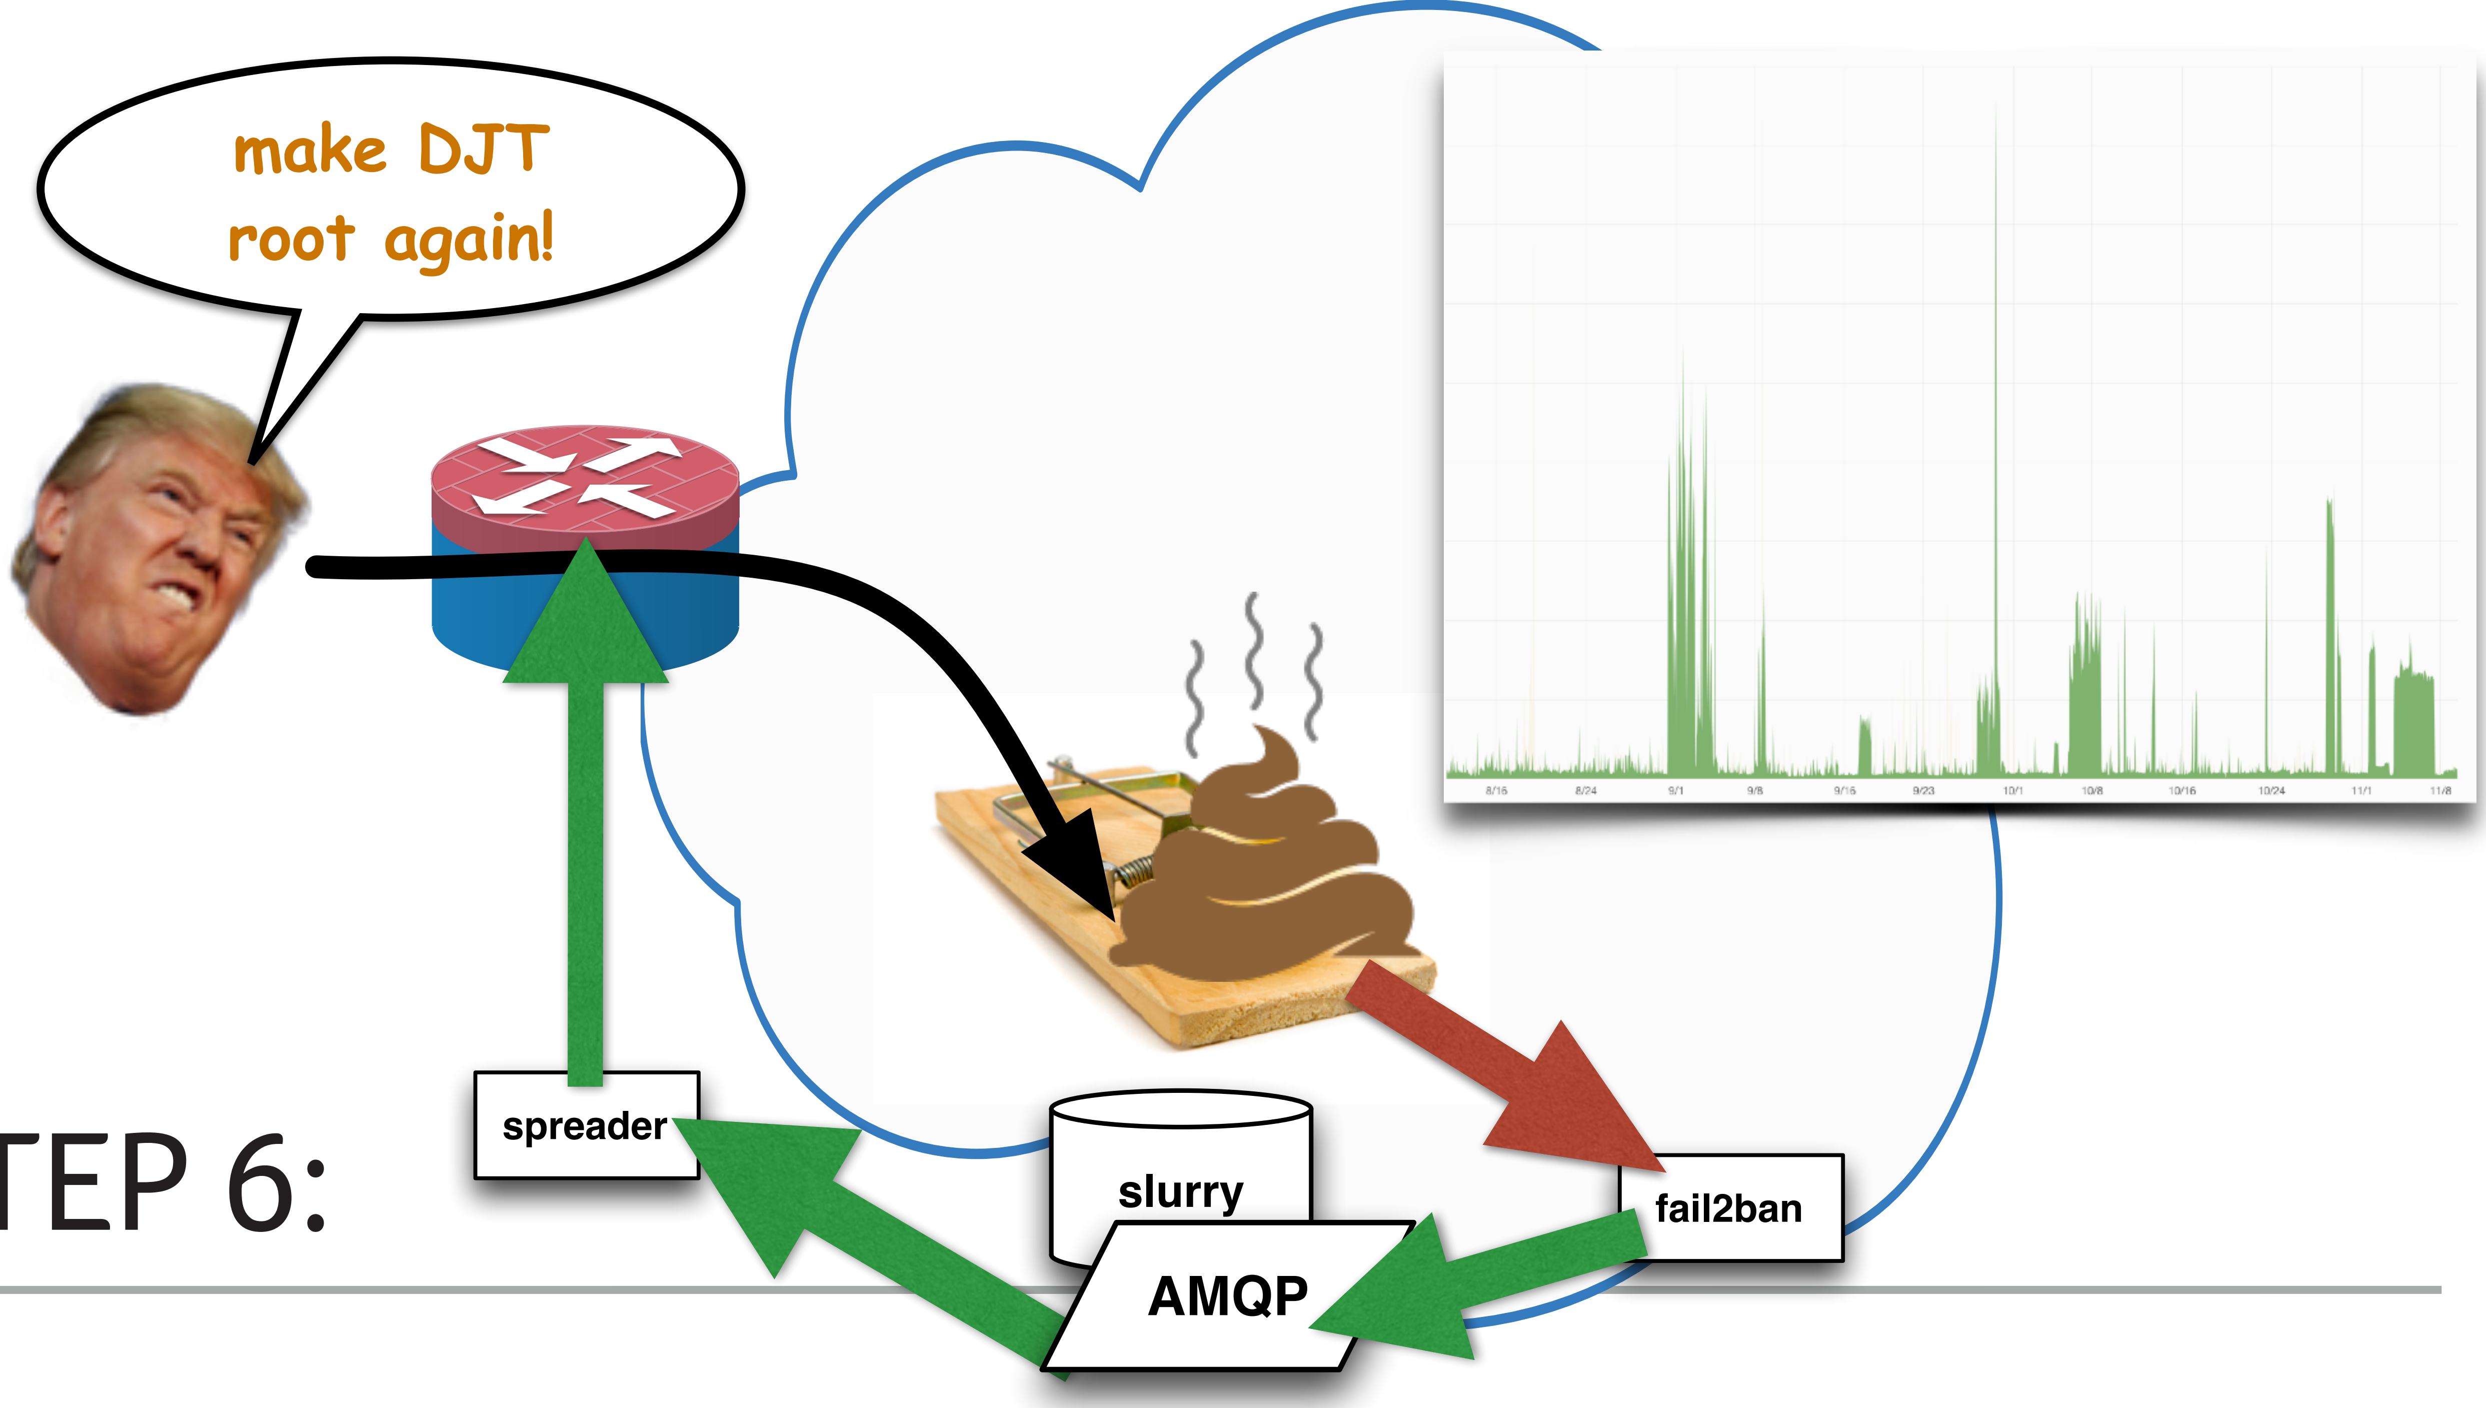

STEP 6:

...AND STAY OUT!

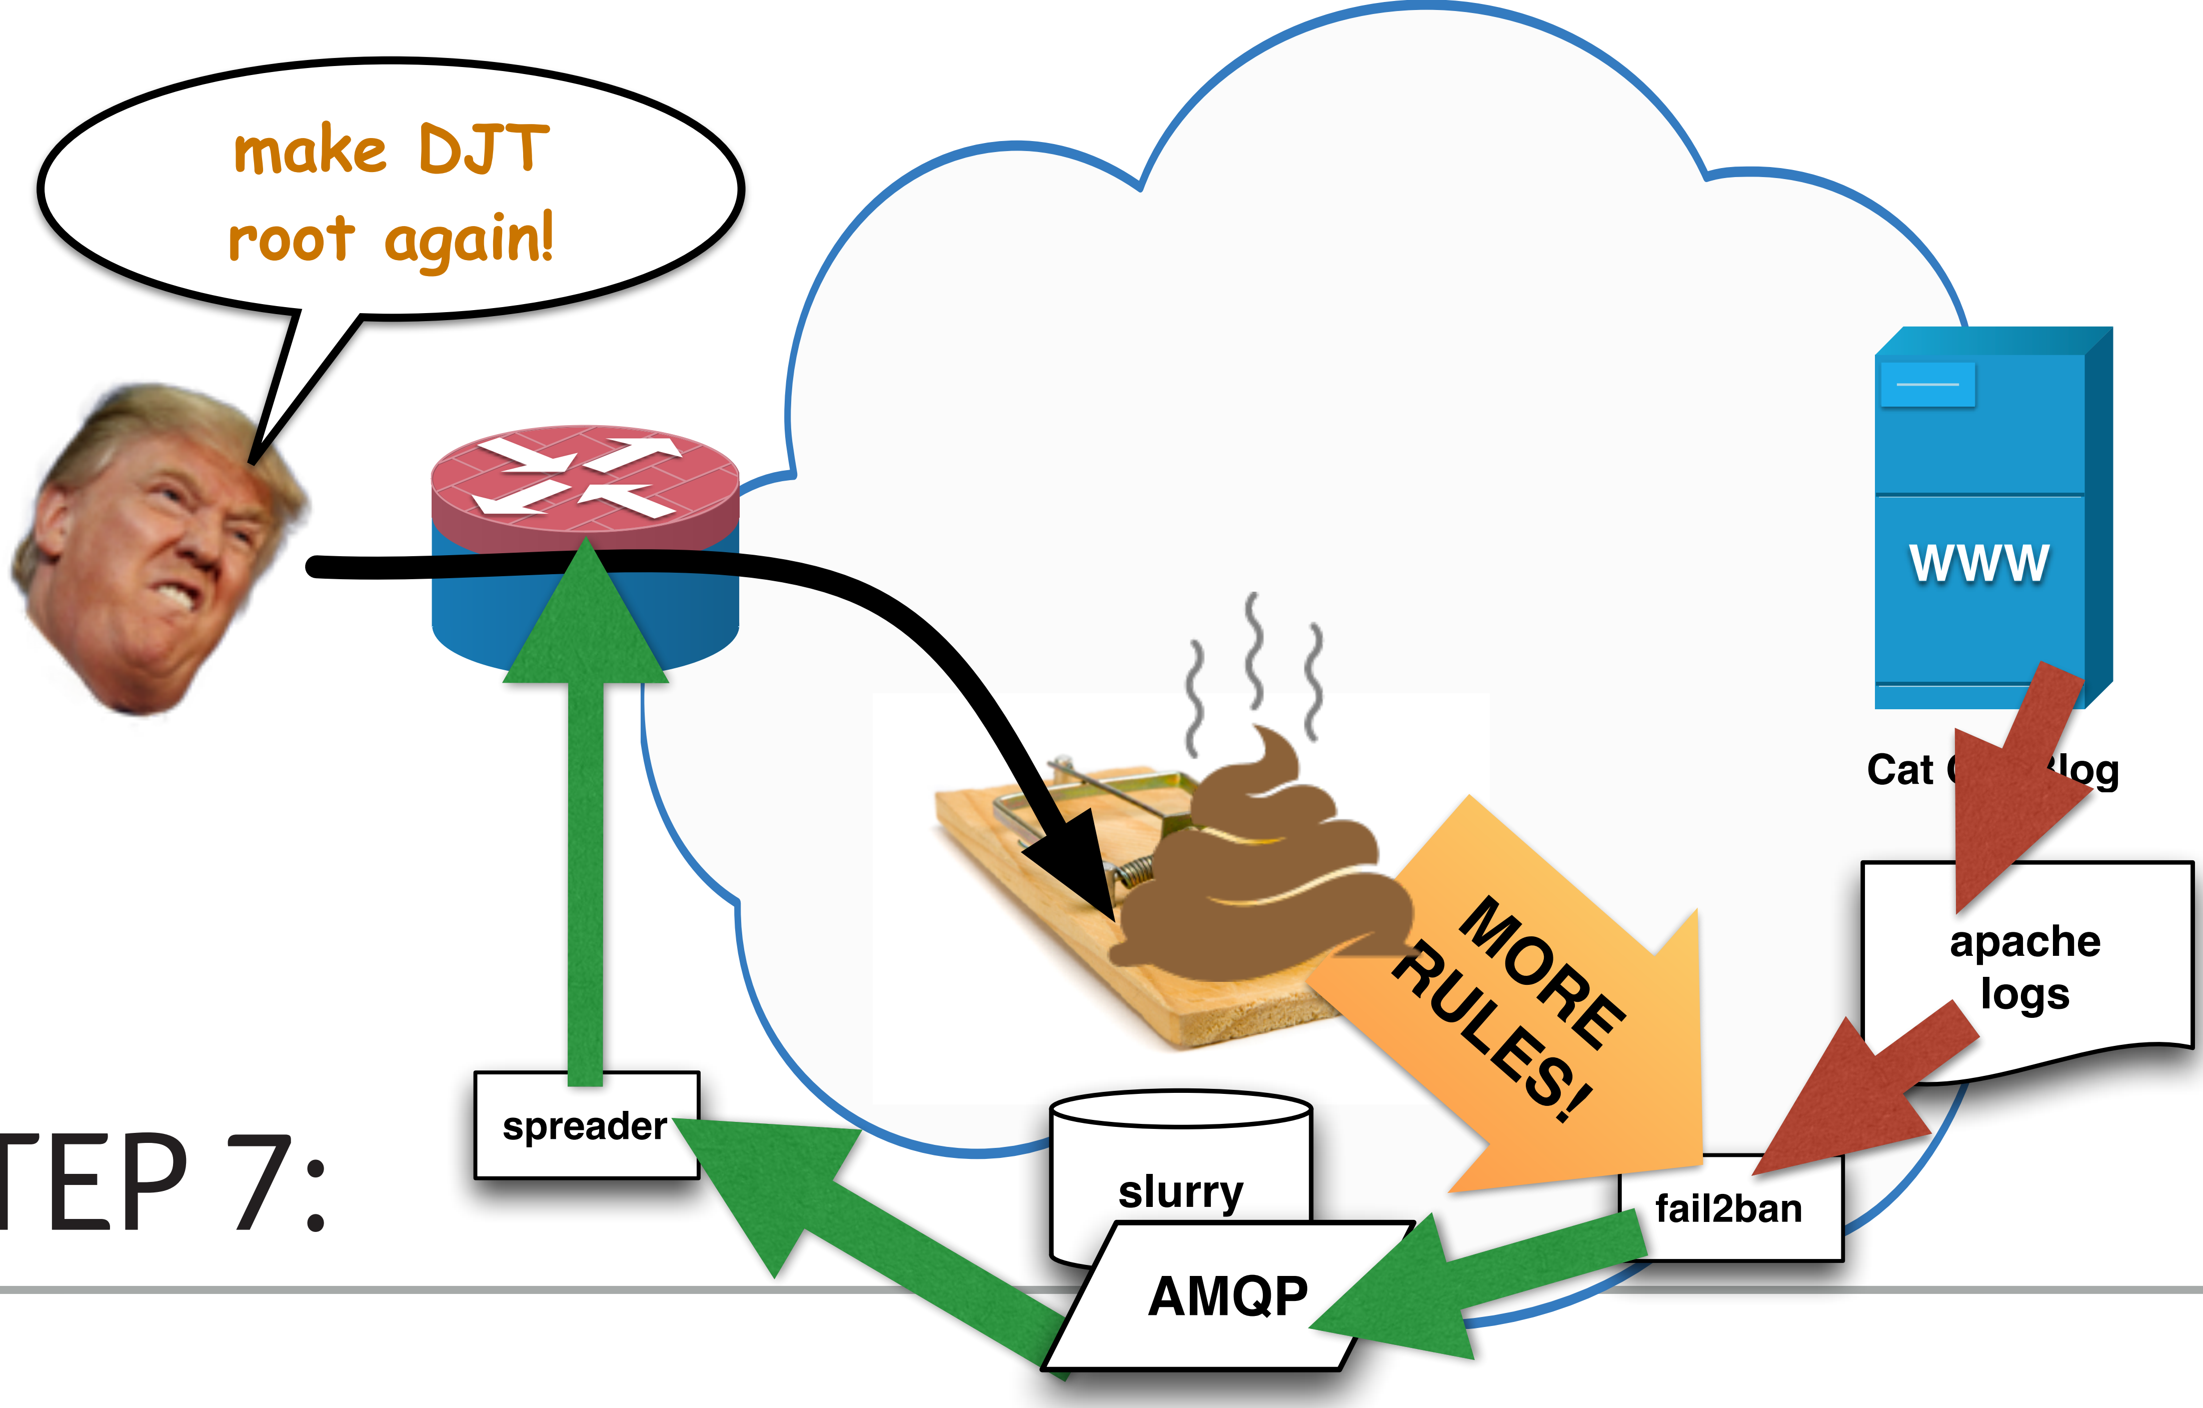

STEP 7:

YOUR NEXT CROP OF LOW-HANGING FRUIT

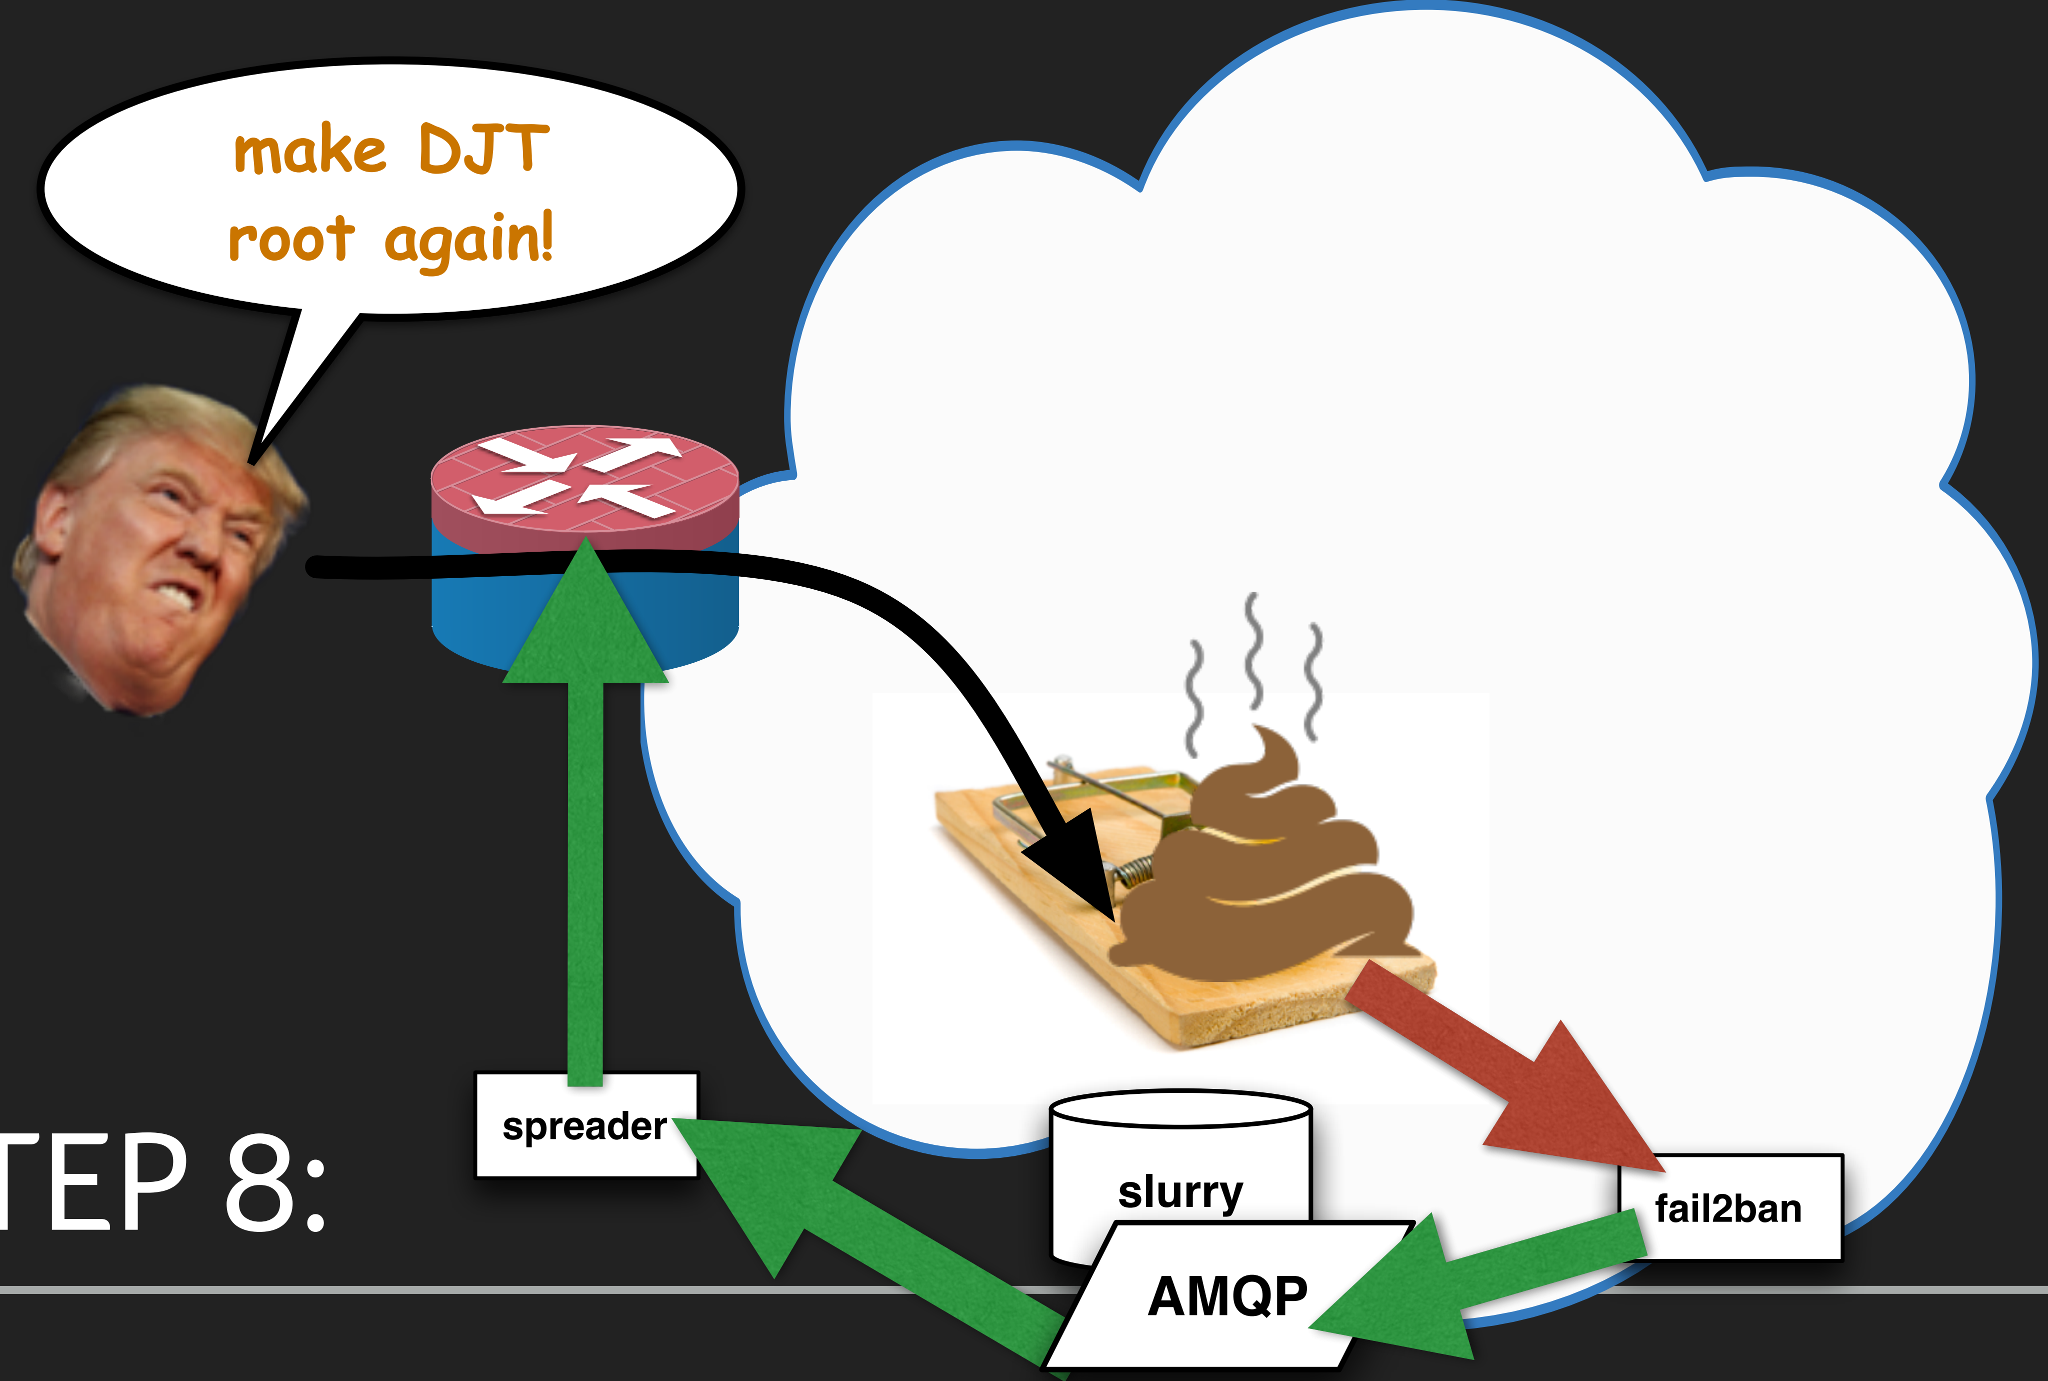

STEP 8:

EXTRA CREDIT

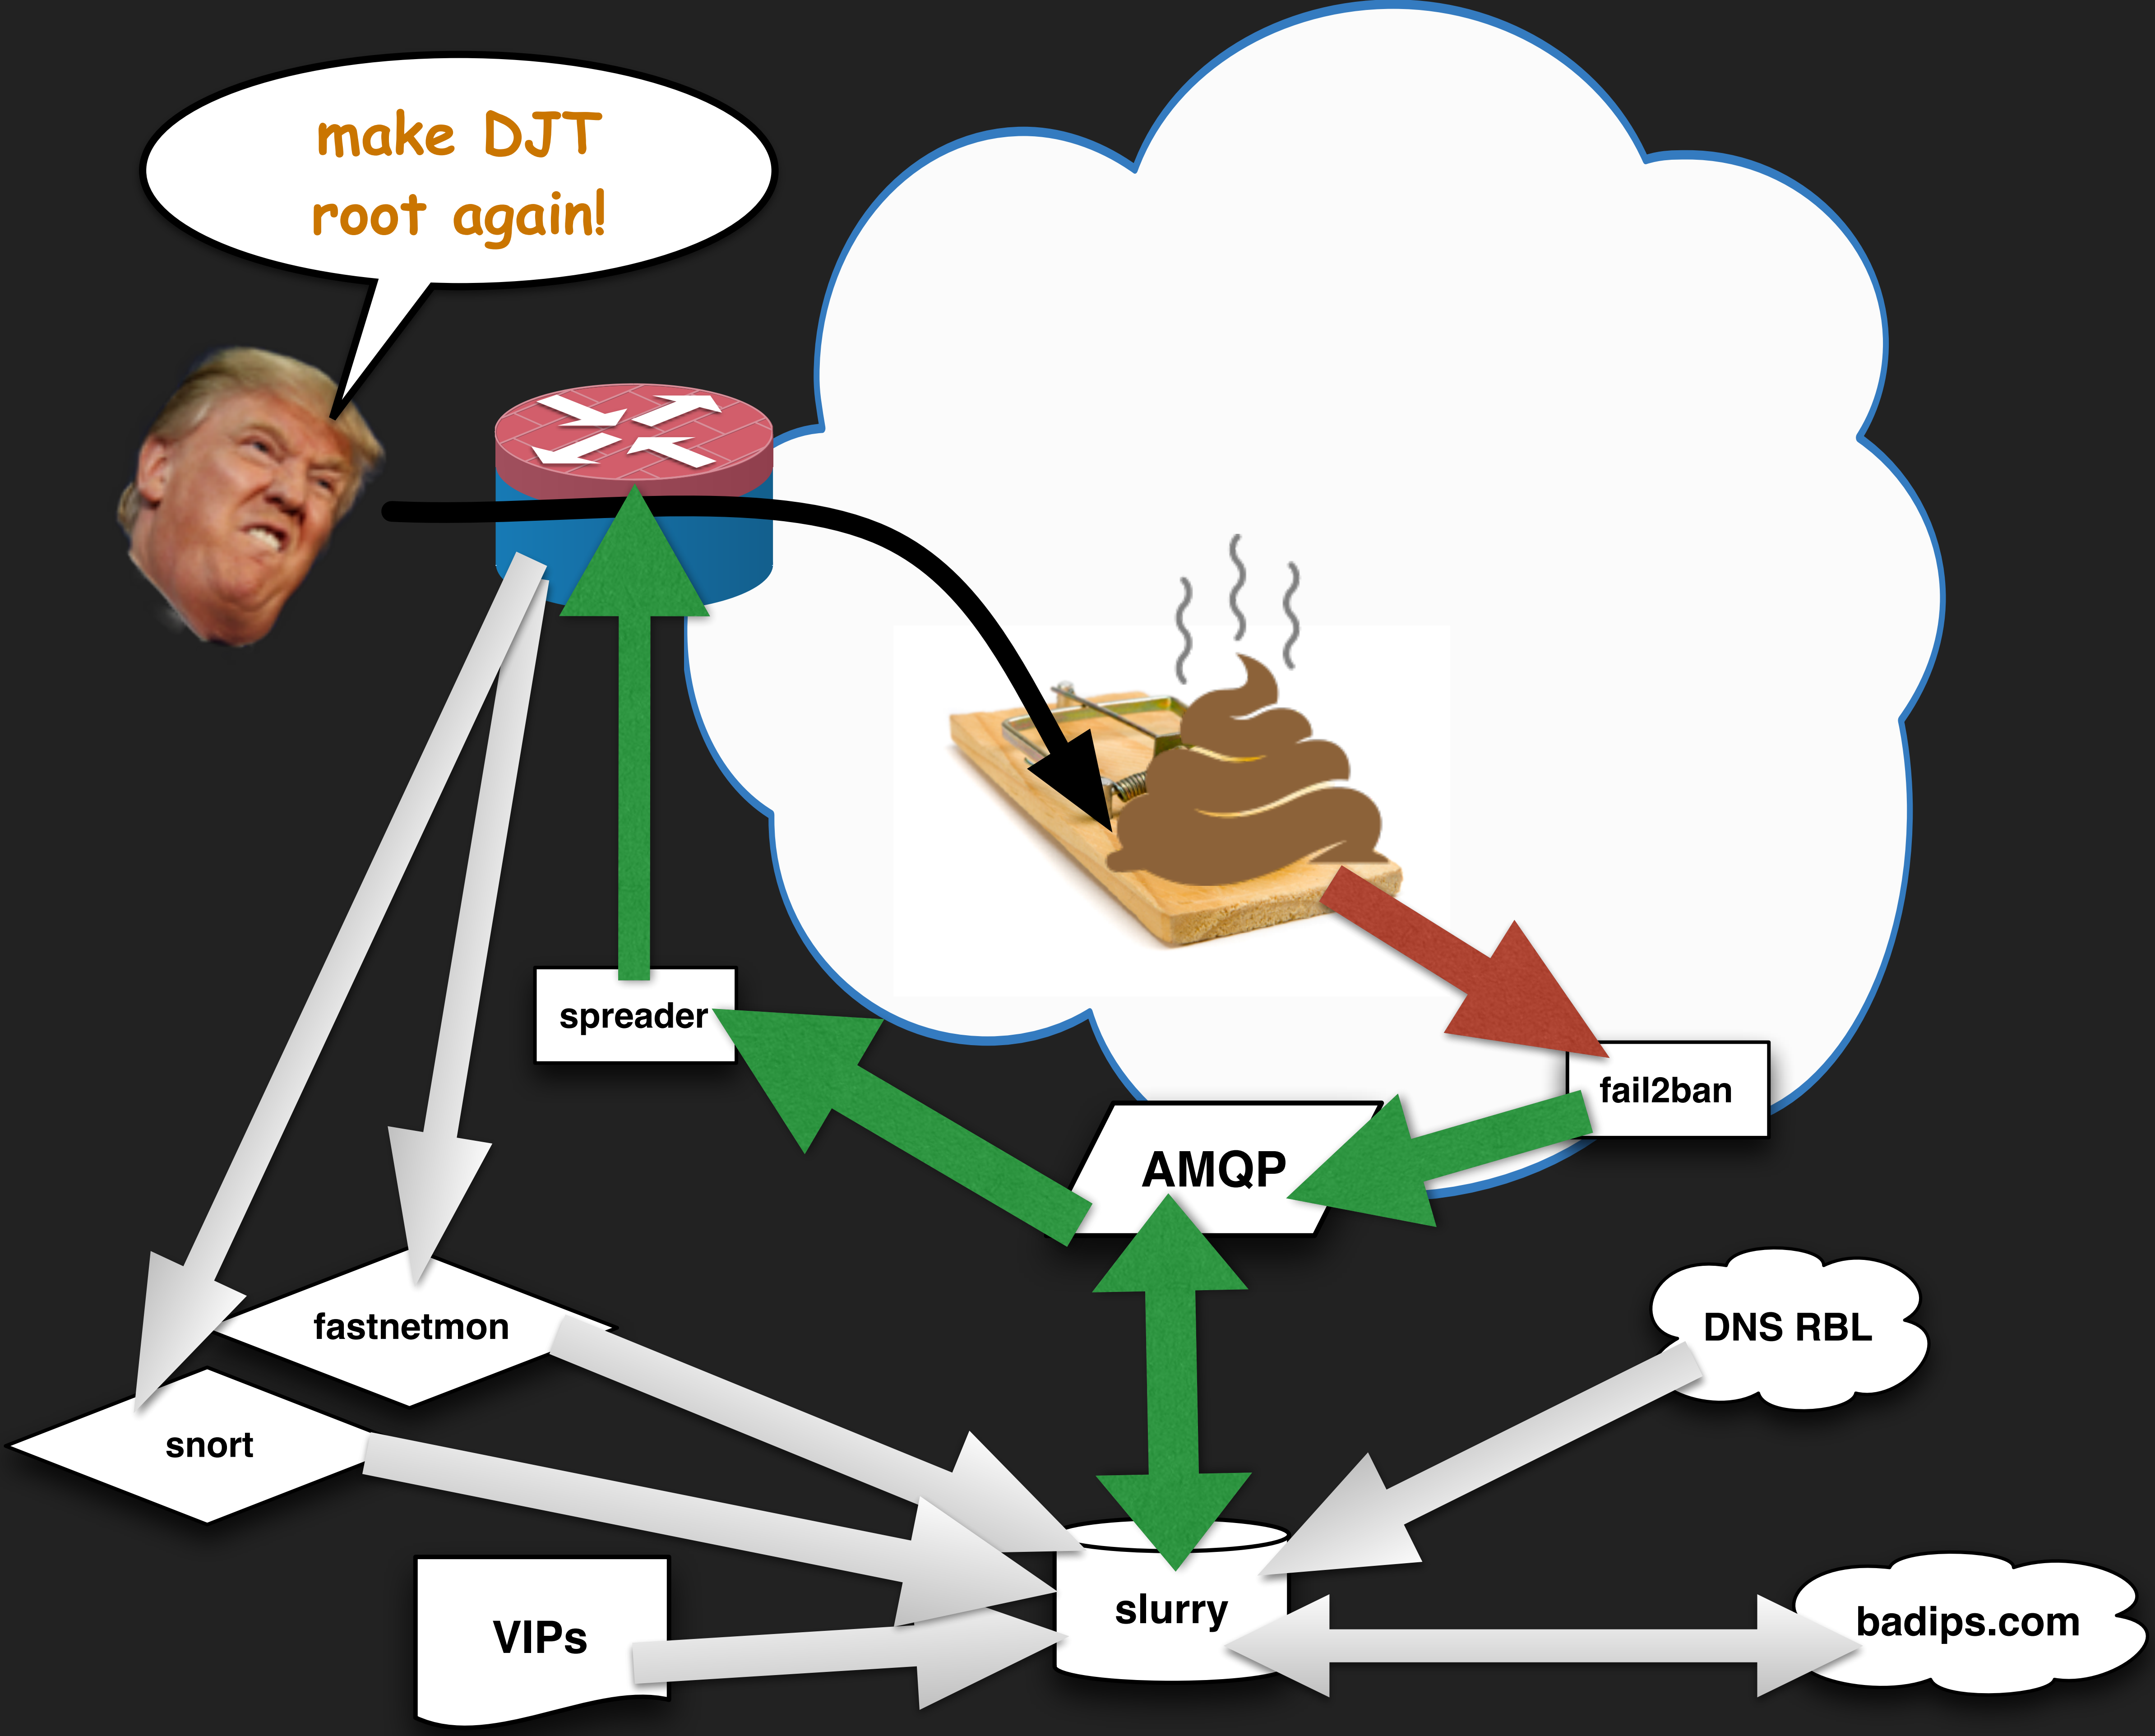

spreader

fail2ban

AMQP

fastnetmon

snort

DNS RBL

VIPs

slurry

badips.com

REFERENCES

- **fail2ban** = tail log files, filter them, perform actions
- **fastnetmon** = DDoS detection with data from **/ip traffic-flow**
- **portsentry** = am I being portscanned?
- **mod_security + OWASP** = Web Application Firewall
- **snort** = intrusion detection system

- GIFs from devopsreactions, securityreactions, honestnetworker

CONCLUSION

-
  ```
  /ip route add gateway=203.0.113.113 routing-mark=shitpit
  ```

-
  ```
  /ip route rule add routing-mark=shitpit table=shitpit
  ```

-
  ```
  /ip firewall mangle add chain=prerouting passthrough=yes
  action=mark-routing new-routing-mark=shitpit
  src-address-list=shitpit
  ```

-
  ```
  /ip firewall address-list add list=shitpit
  address=192.0.2.69/32 timeout=1m
  ```

THANKS FOR LISTENING!

ANY QUESTIONS?

e: marek@faelix.net
t: @maznu
w: https://faelix.net/