swoop

Did You Drop Something?

Fast carrier DDoS detection & mitigation, at scale.

swoop.com.au

This is the story of how the Swoop IP Transit network evolved over the past 7 years to deal with emerging and evolving DDoS threats.

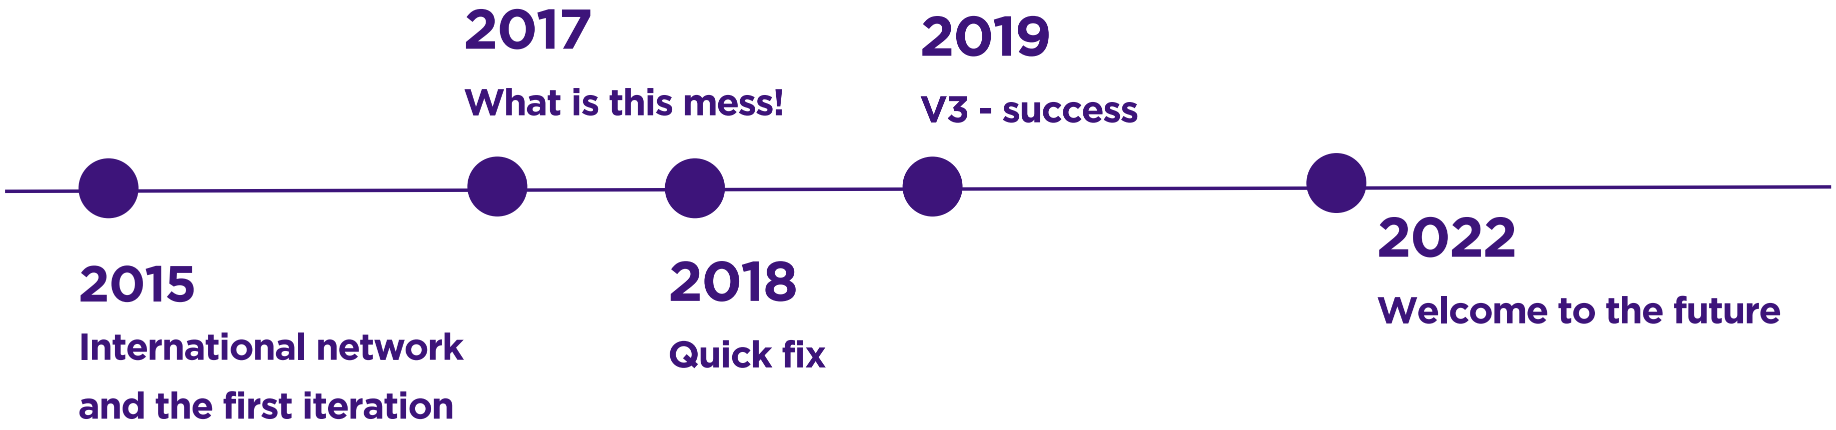

2015
International network
and the first iteration

2017
What is this mess!

2018
Quick fix

2019
V3 - success

2022
Welcome to the future

What is a DDoS?

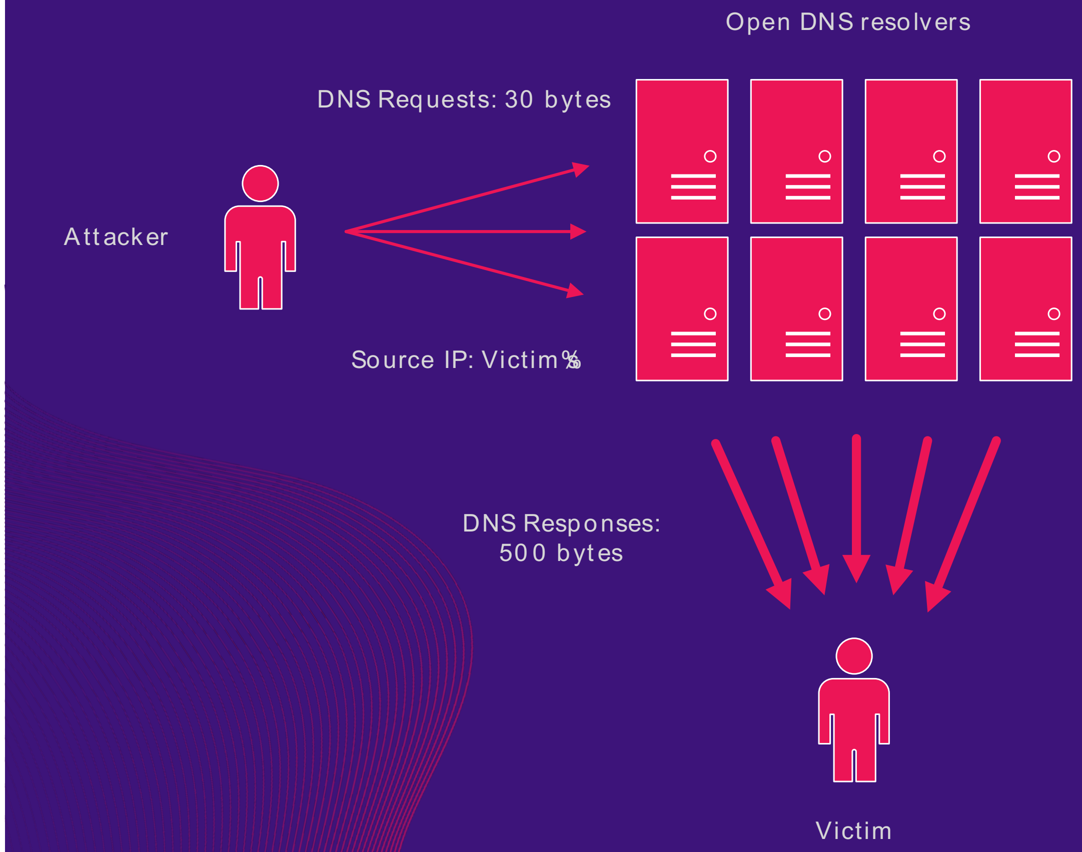

Open DNS resolvers

DNS Requests: 30 bytes

Attacker

Source IP: Victim%s

DNS Responses:
500 bytes

Victim

We probably all know the basics...

What can we as a service provider (the innocent middleman), hope to do about this?

Protect our network and customers

Stop the pipes being overwhelmed

Use our budget wisely

The budget: $0.00

React quickly

Less than 10 seconds to detect and mitigate

The case for "roll-your-own"

Inline hardware solutions have drawbacks…

…and we had a guy that writes code

swoop

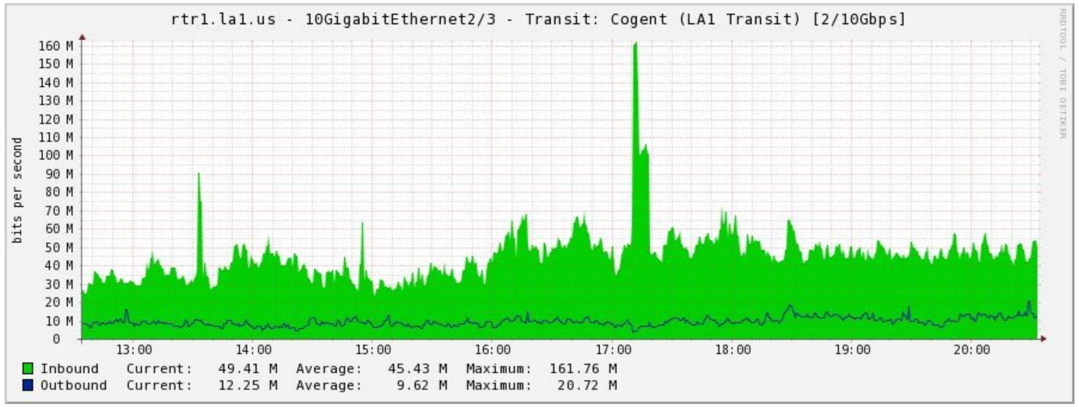

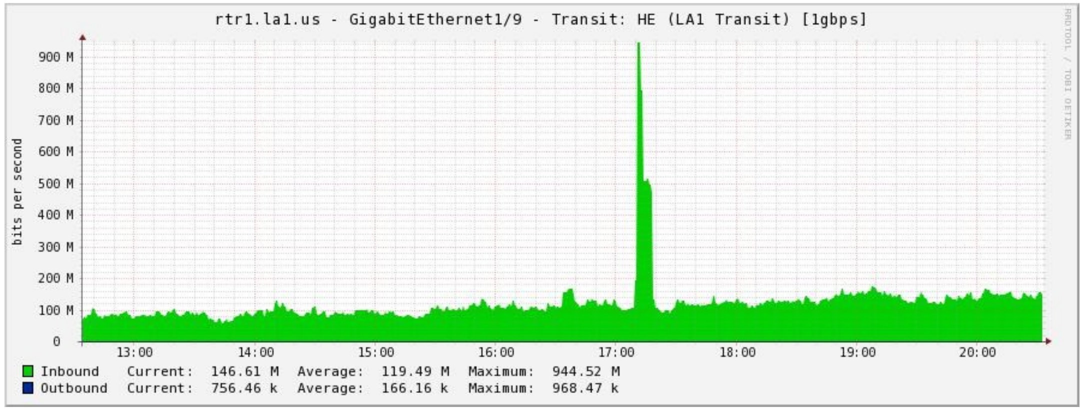

The year was 2015…

We had just begun to expand globally, with our first international PoPs established in CoreSite LA1 and Telehouse in London

The first iteration:
“Hammer”

- Written in NodeJS
- Brocade CER-RT based IP network limited to only sFlow
- Low Network Edge Capacity

Mitigation Approach

- Forge (Hijack) a new route in BGP for the /24 of the attacked /32
- Set NEXT_HOP to the original prefix's NEXT_HOP
- Tag it with communities that both cause it to be announced to the DDoS protection provider, and *not* announced to any of our transit or peering

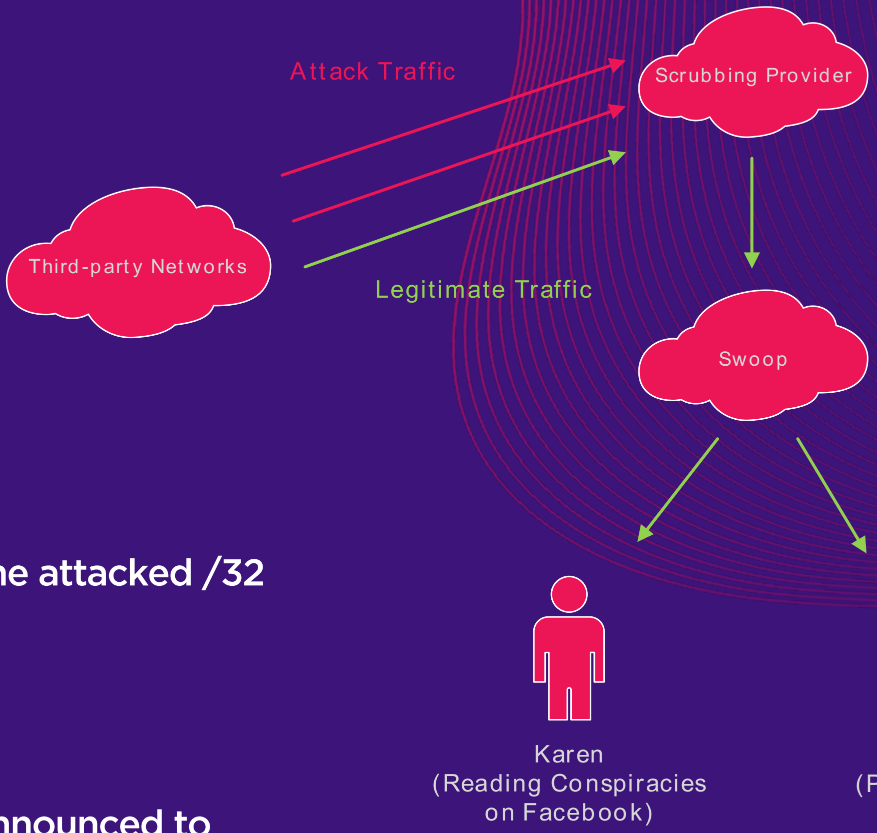

Don’t do this.

- Route hijacking, even for legit reasons, is bad.
- Relying on a third party for scrubbing is bad.
- Stuck and stale routes… are bad.
- Rewriting of origin ASN is bad.
- Static next hops are, you guessed it, bad.

a quick fix-
Kentik

- **Fortunately, at the time, we used Kentik for their awesome traffic analysis, but they also offered an “alerting and actions” architecture for DDoS detection and mitigation**

- **Custom policies allowed us to define thresholds, e.g. for common DRDoS source ports to a single IP address exceeding 50Mbps, and then take an action (at the time, only blackholing was available)**

- **While this was a HUGE improvement on our 5-10 minutes, we still wanted faster**

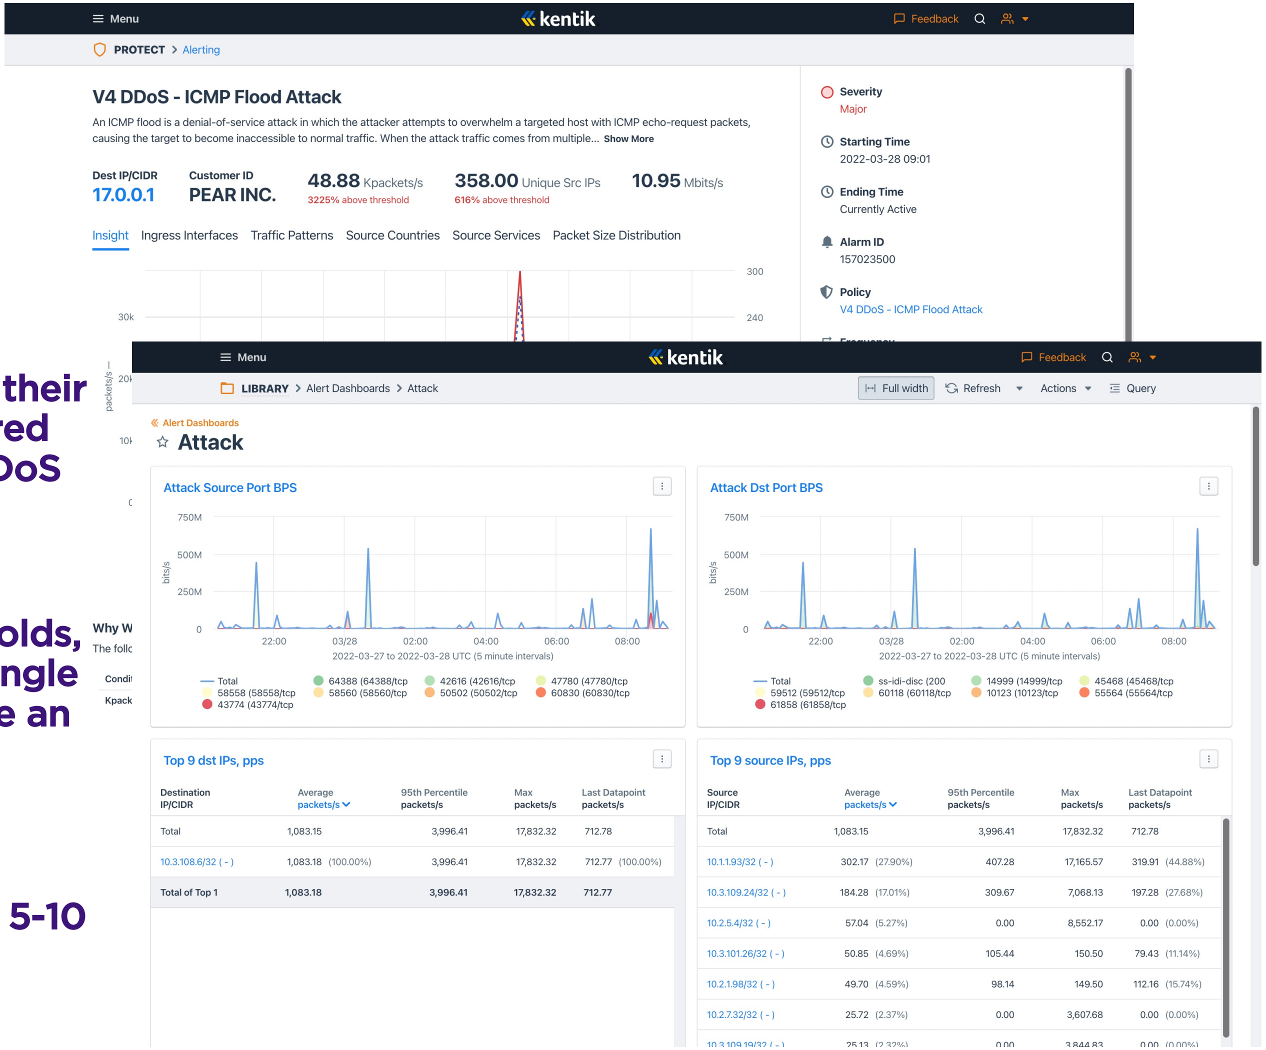

V4 DDoS - ICMP Flood Attack

An ICMP flood is a denial-of-service attack in which the attacker attempts to overwhelm a targeted host with ICMP echo-request packets, causing the target to become inaccessible to normal traffic. When the attack traffic comes from multiple... Show More

Dest IP/CIDR: 17.0.0.1
Customer ID: PEAR INC.
48.88 Kpackets/s — 3225% above threshold
358.00 Unique Src IPs — 616% above threshold
10.95 Mbits/s

Severity: Major
Starting Time: 2022-03-28 09:01
Ending Time: Currently Active
Alarm ID: 157023500
Policy: V4 DDoS - ICMP Flood Attack

Attack

Top 9 dst IPs, pps

Destination IP/CIDR	Average packets/s	95th Percentile packets/s	Max packets/s	Last Datapoint packets/s
Total	1,083.15	3,996.41	17,832.32	712.78
10.3.108.6/32 (-)	1,083.18 (100.00%)	3,996.41	17,832.32	712.77 (100.00%)
Total of Top 1	**1,083.18**	**3,996.41**	**17,832.32**	**712.77**

Top 9 source IPs, pps

Source IP/CIDR	Average packets/s	95th Percentile packets/s	Max packets/s	Last Datapoint packets/s
Total	1,083.15	3,996.41	17,832.32	712.78
10.1.1.93/32 (-)	302.17 (27.90%)	407.28	17,165.57	319.91 (44.88%)
10.3.109.24/32 (-)	184.28 (17.01%)	309.67	7,068.13	197.28 (27.68%)
10.2.5.4/32 (-)	57.04 (5.27%)	0.00	8,552.17	0.00 (0.00%)
10.3.101.26/32 (-)	50.85 (4.69%)	105.44	150.50	79.43 (11.14%)
10.2.1.98/32 (-)	49.70 (4.59%)	98.14	149.50	112.16 (15.74%)
10.2.7.32/32 (-)	25.72 (2.37%)	0.00	3,607.68	0.00 (0.00%)
10.3.109.19/32 (-)	25.13 (2.32%)	0.00	3,844.83	0.00 (0.00%)

2019

(remember no COVID?)

Edge capacity for days

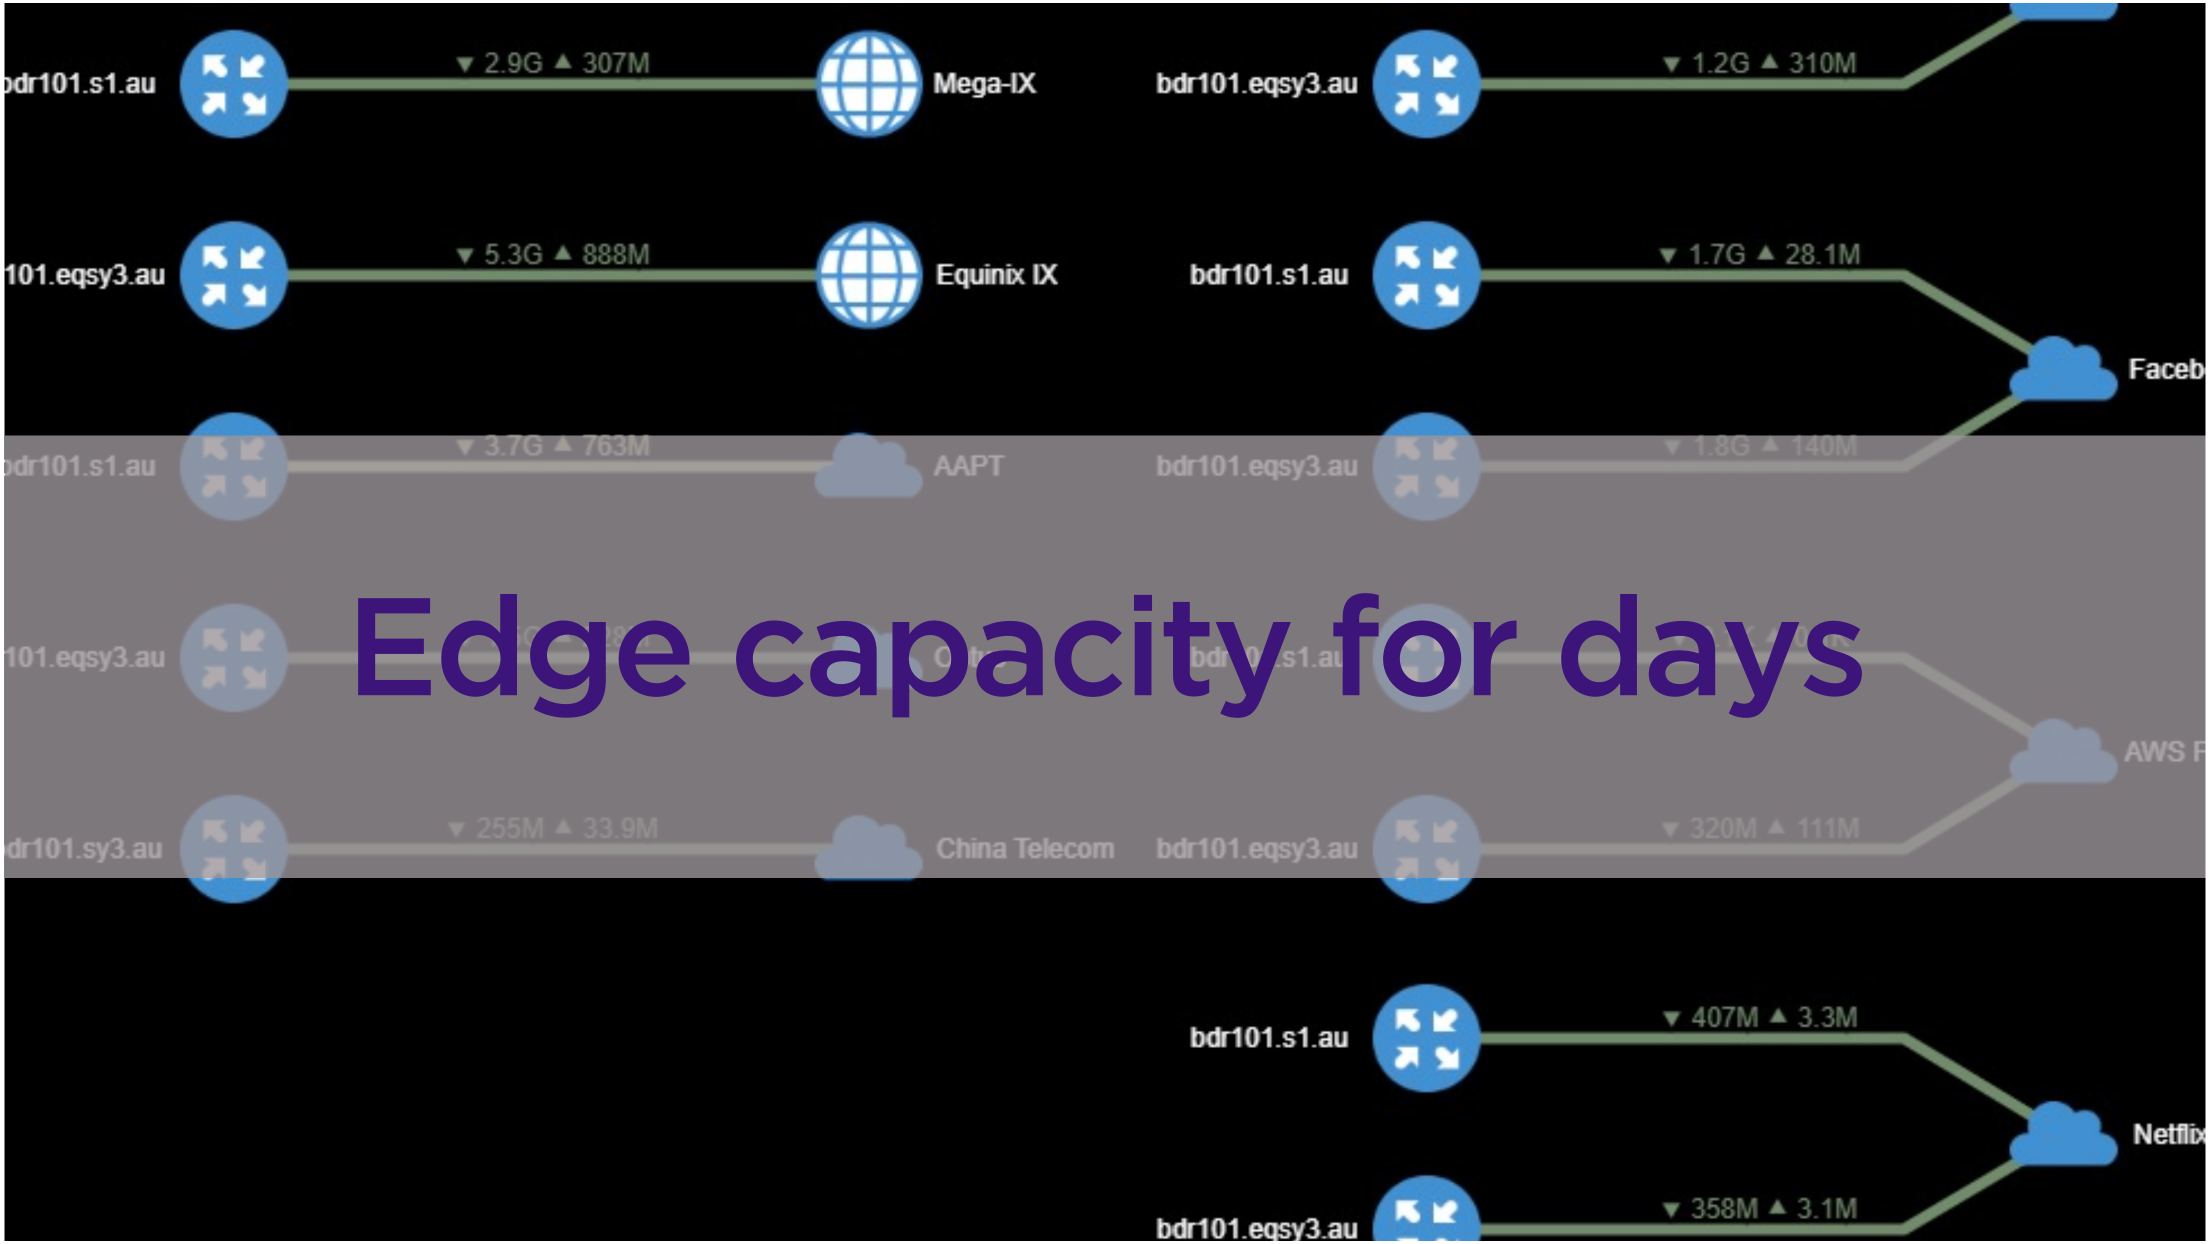

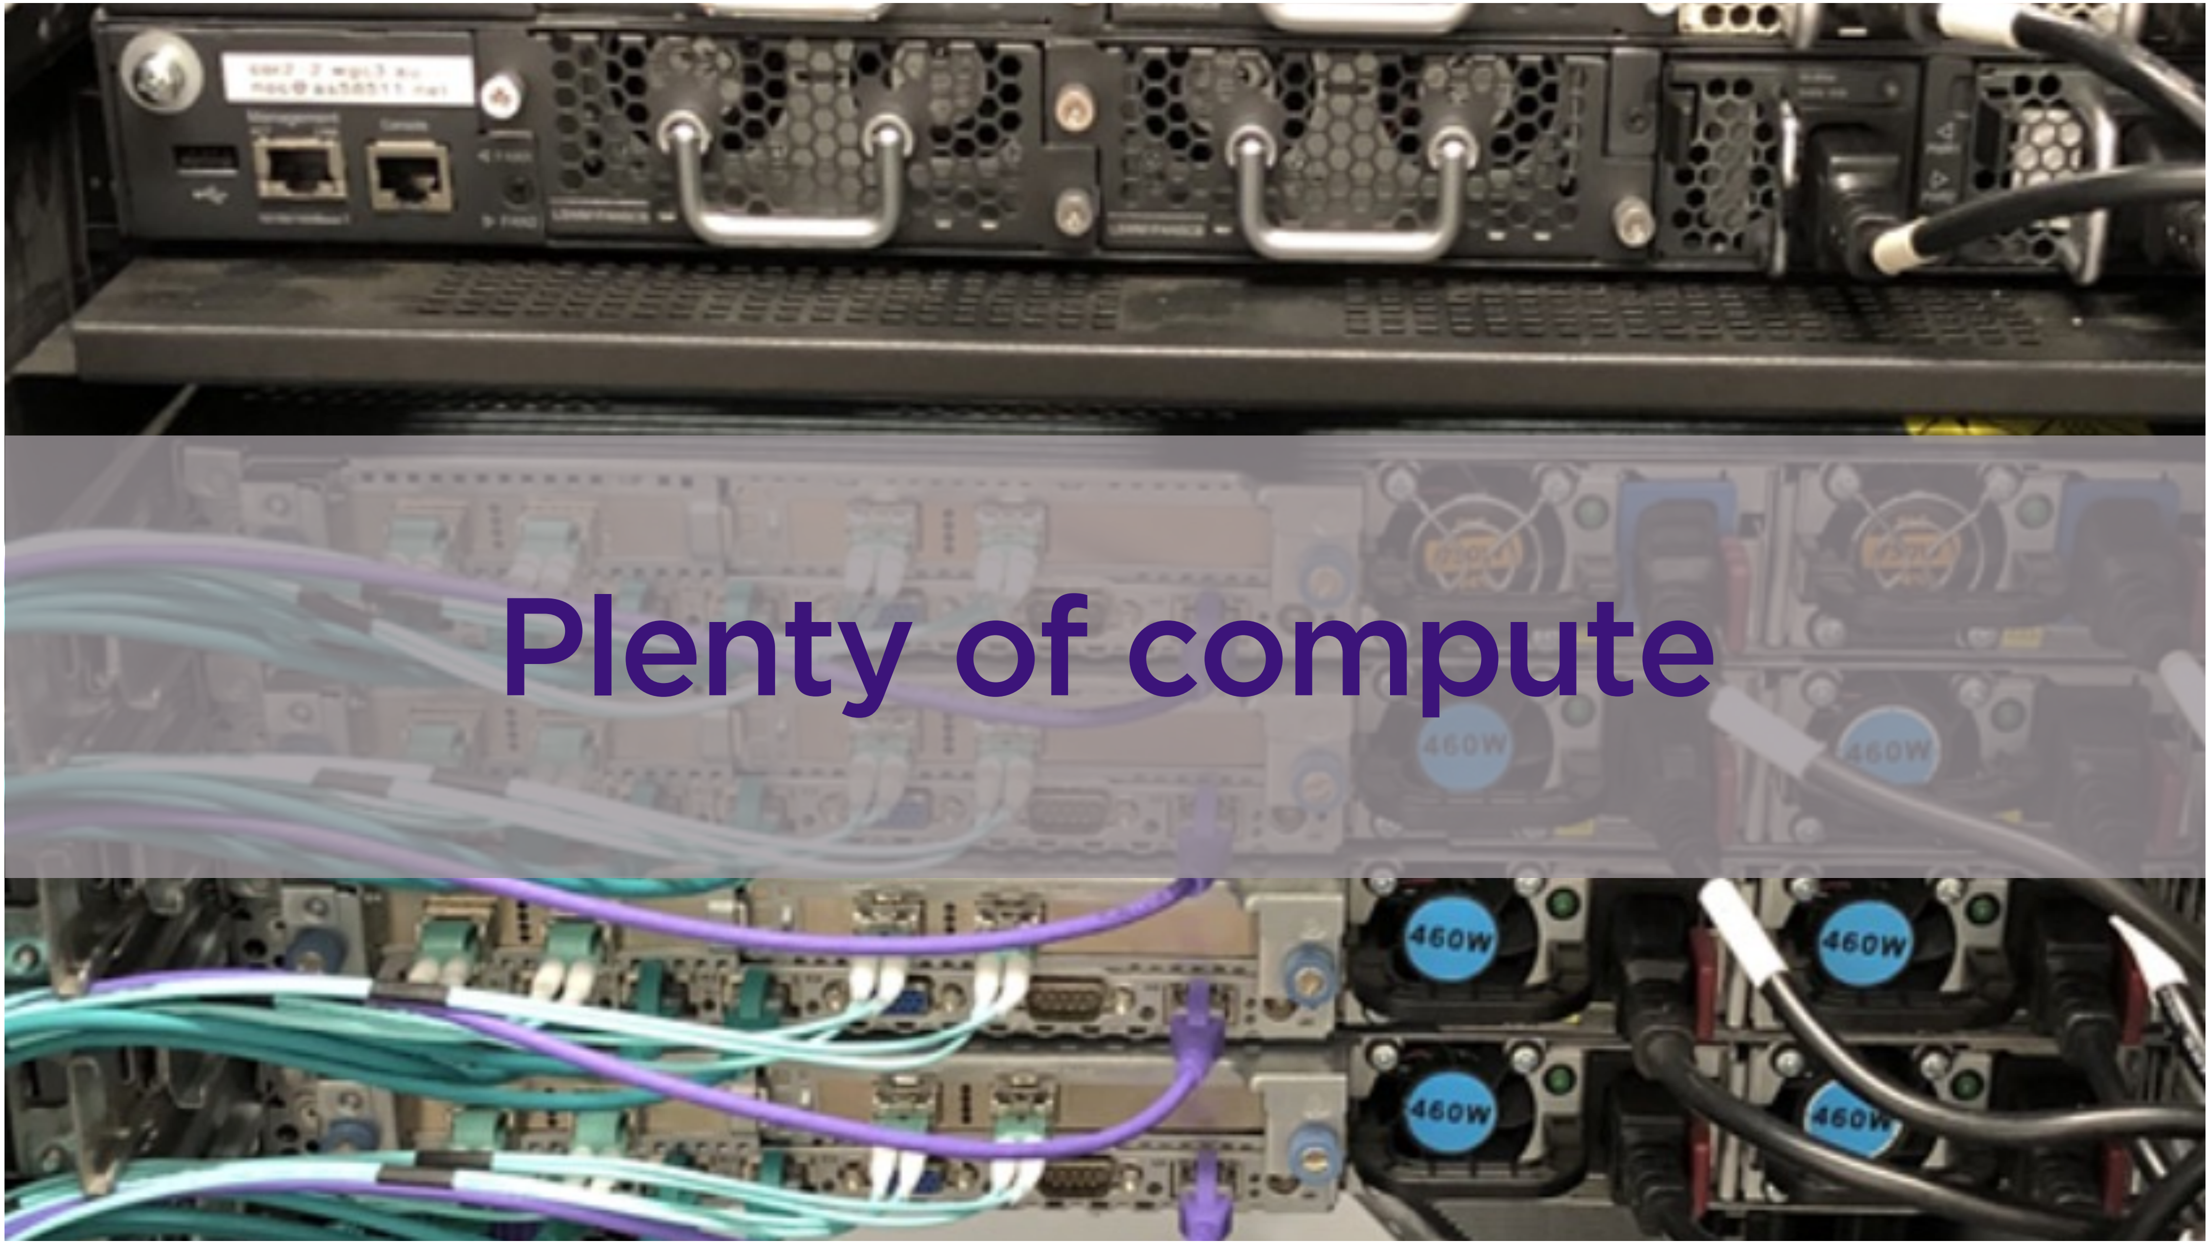

Plenty of compute

A guy that writes code

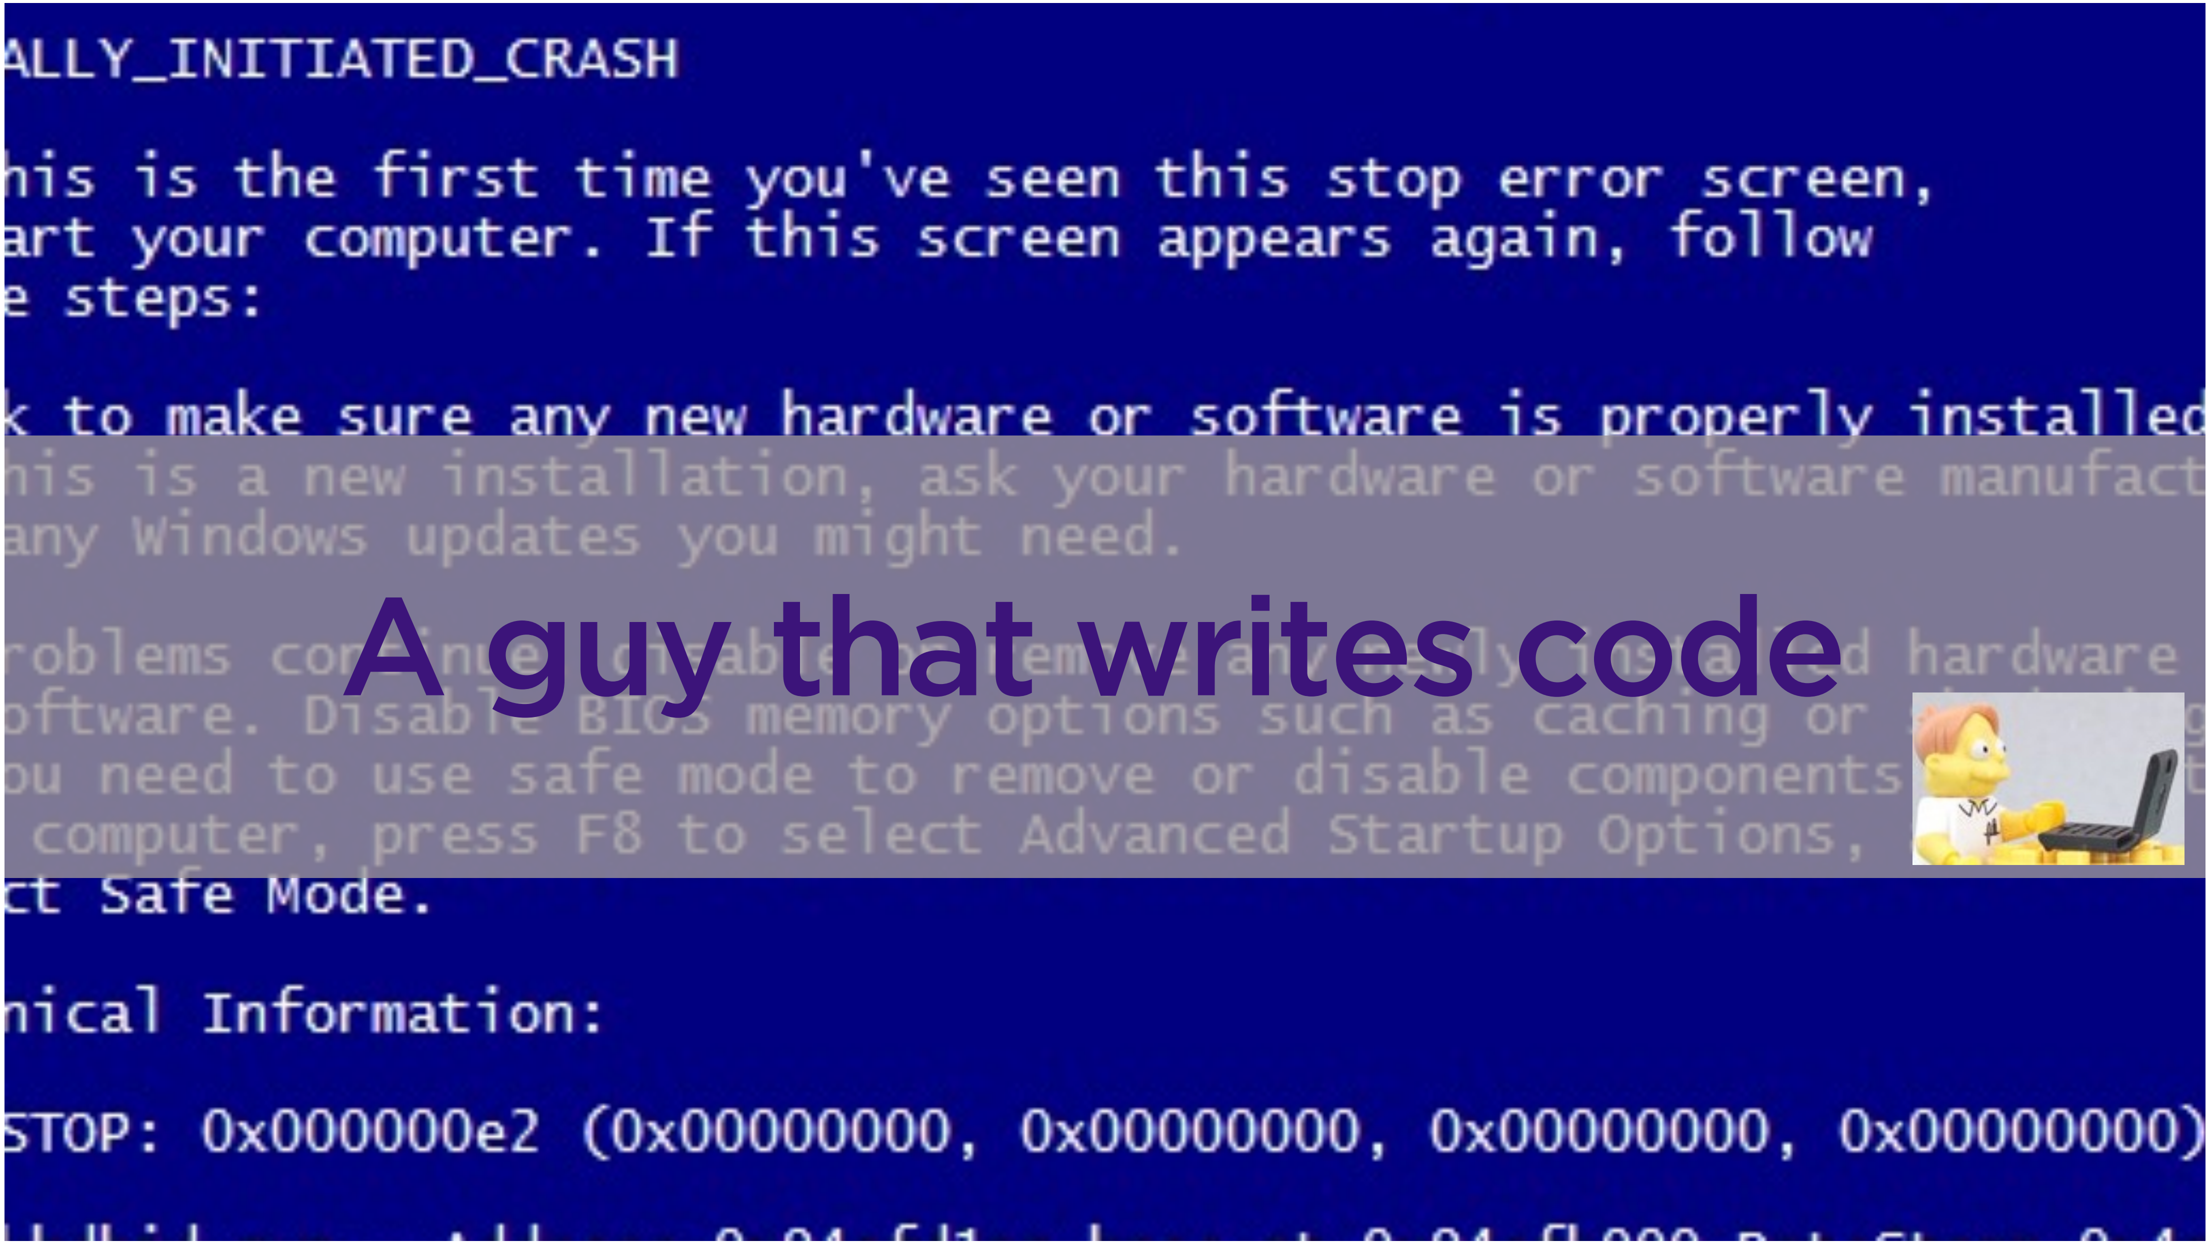

```
ALLY_INITIATED_CRASH

his is the first time you've seen this stop error screen,
art your computer. If this screen appears again, follow
e steps:

k to make sure any new hardware or software is properly installed
his is a new installation, ask your hardware or software manufact
any Windows updates you might need.

roblems continue, disable or remove any newly installed hardware
oftware. Disable BIOS memory options such as caching or shadowing
ou need to use safe mode to remove or disable components, t
 computer, press F8 to select Advanced Startup Options,
ct Safe Mode.

nical Information:

STOP: 0x000000e2 (0x00000000, 0x00000000, 0x00000000, 0x00000000)
```

The third iteration:
Sentinel

What are we trying to solve?

Improvements on the previous approaches

No Third-Party Scrubbing / No Blackholing

Improve Scalability & Improve Response Time

Use all of our Budget

(The budget: ...still $0.00)

Sentinel

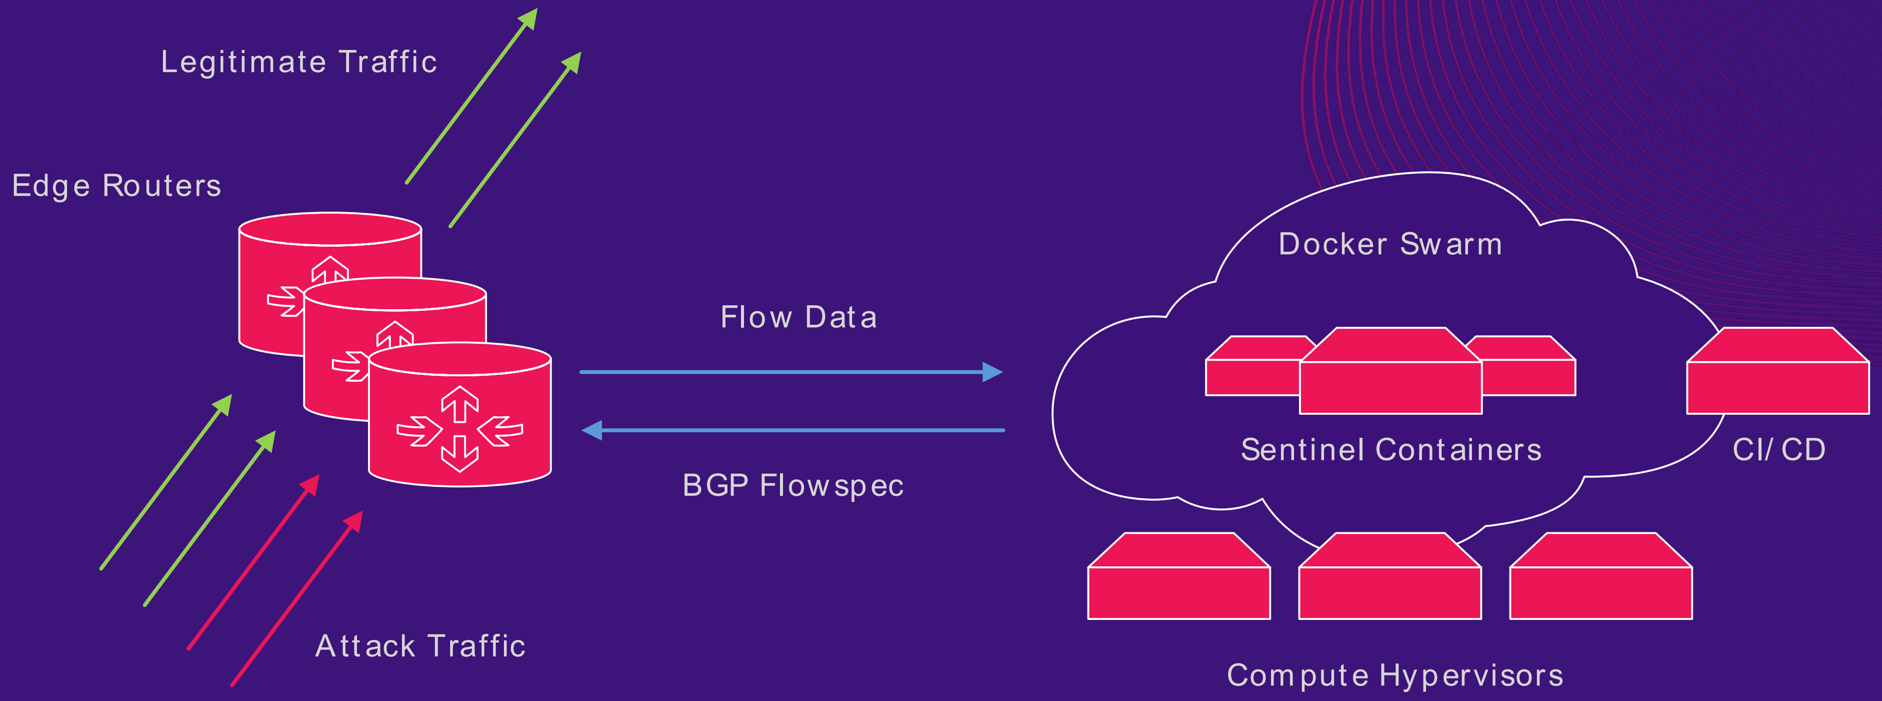

swoop

Infrastructure as Code

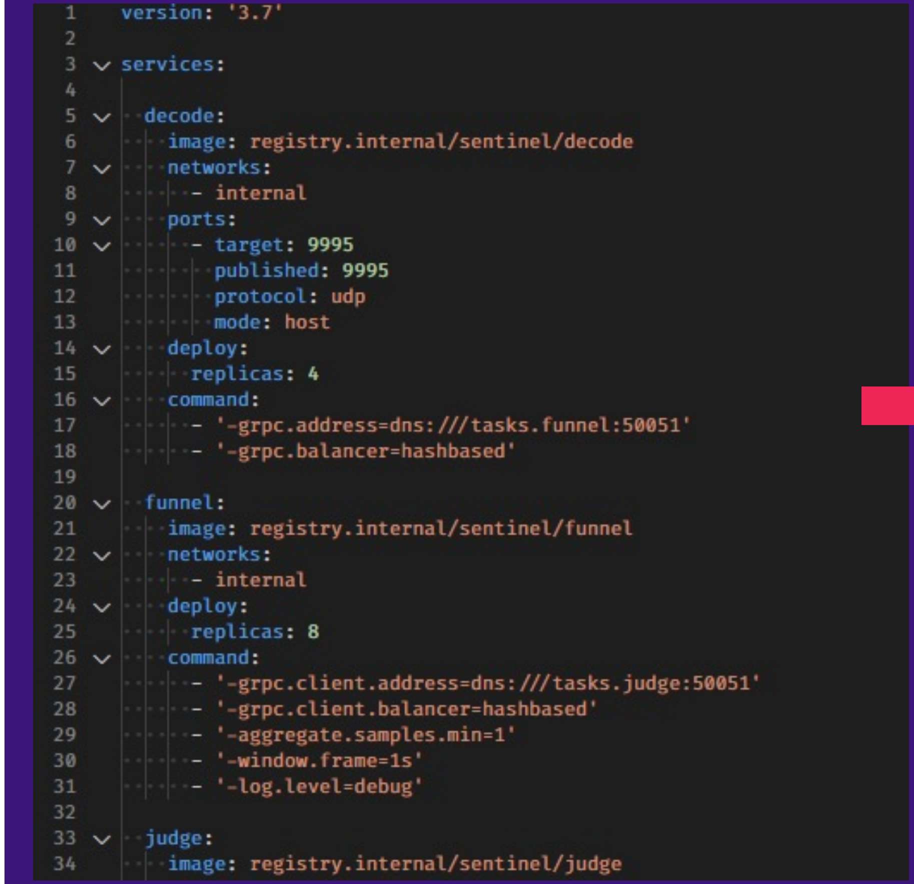

```yaml
version: '3.7'

services:

  decode:
    image: registry.internal/sentinel/decode
    networks:
      - internal
    ports:
      - target: 9995
        published: 9995
        protocol: udp
        mode: host
    deploy:
      replicas: 4
    command:
      - '-grpc.address=dns:///tasks.funnel:50051'
      - '-grpc.balancer=hashbased'

  funnel:
    image: registry.internal/sentinel/funnel
    networks:
      - internal
    deploy:
      replicas: 8
    command:
      - '-grpc.client.address=dns:///tasks.judge:50051'
      - '-grpc.client.balancer=hashbased'
      - '-aggregate.samples.min=1'
      - '-window.frame=1s'
      - '-log.level=debug'

  judge:
    image: registry.internal/sentinel/judge
```

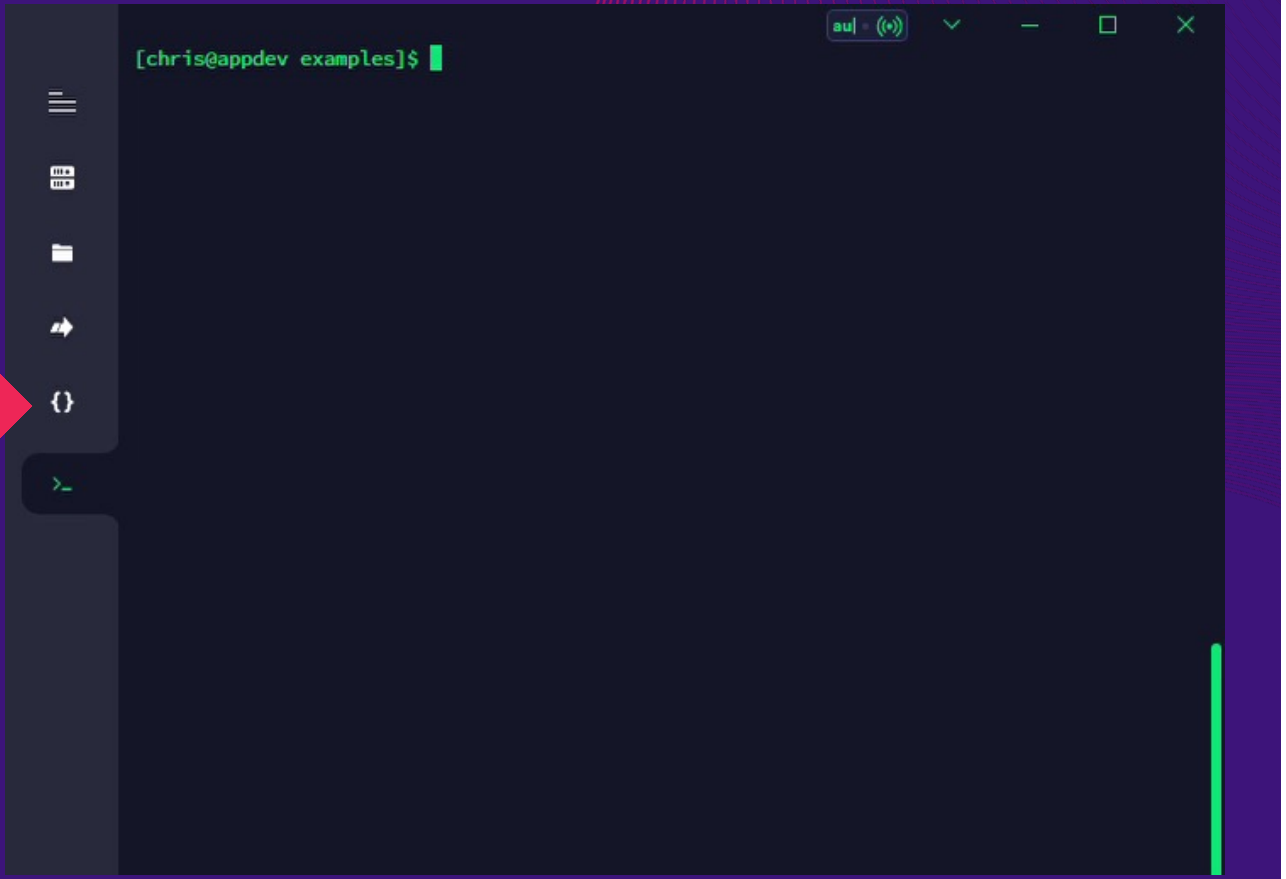

```
[chris@appdev examples]$ 
```

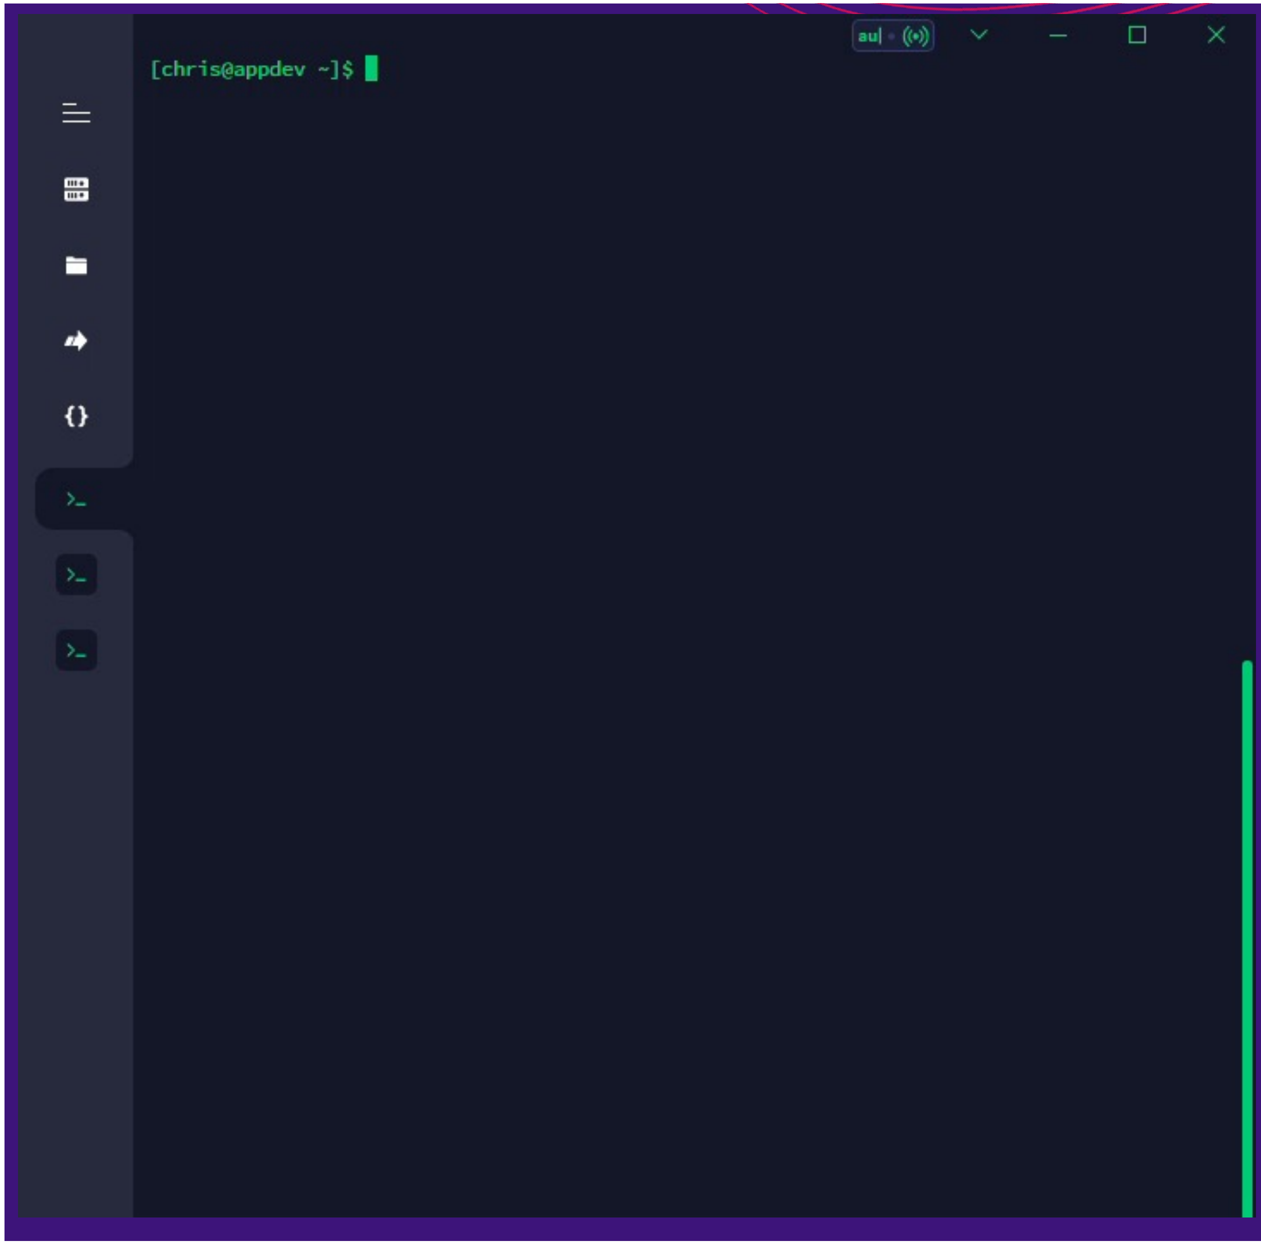

```
[chris@appdev ~]$ 
```

Scaling!

(or, how to burn CPU time)

swoop

How fast is it?

World-wide mitigation in under 10 seconds

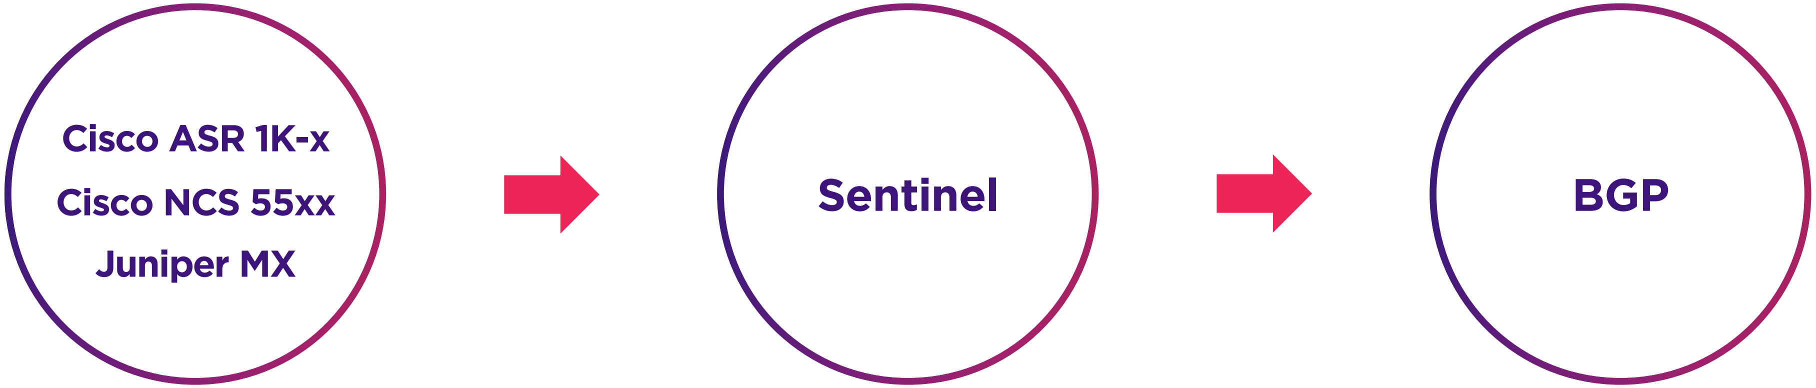

Cisco ASR 1K-x

Cisco NCS 55xx

Juniper MX

Sentinel

BGP

NetFlow: 0-10s

Routers gather flow state and export to collector

Evaluate: 0-3s

Compare flow data to policies.

Mitigate: 0-2s

Export FlowSpec rules into BGP, routers install them and start dropping the rude packets

Flexible policies

Policies allow us to define thresholds for specific traffic classes, or attack types

LDAP

SSDP

DNS

SNMP

NTP

Known DRDoS
< 50mbps

High PPS

Generic TCP, UDP
(and VPN)

BGP Flowspec

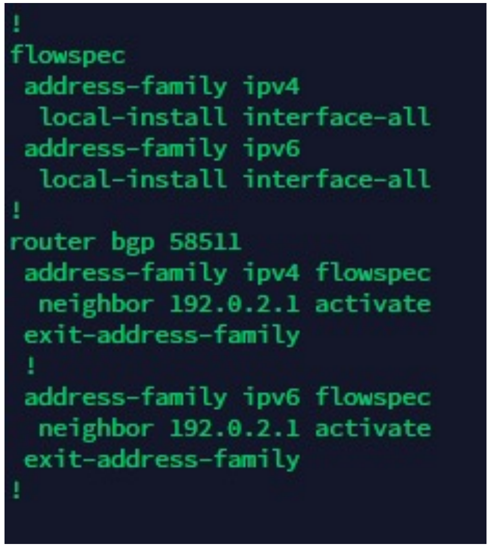

```
!
flowspec
 address-family ipv4
  local-install interface-all
 address-family ipv6
  local-install interface-all
!
router bgp 58511
 address-family ipv4 flowspec
  neighbor 192.0.2.1 activate
 exit-address-family
  !
 address-family ipv6 flowspec
  neighbor 192.0.2.1 activate
 exit-address-family
!
```

Flowspec is a BGP extension that allows traffic rules (match criteria + action) to be propagated to devices via BGP

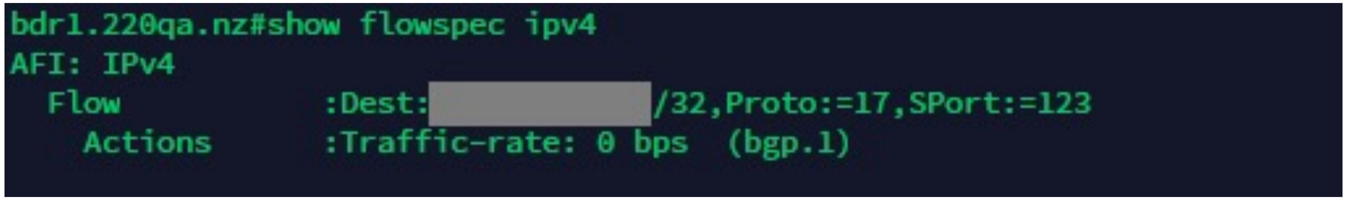

```
bdr1.220qa.nz#show flowspec ipv4
AFI: IPv4
  Flow           :Dest:[redacted]/32,Proto:=17,SPort:=123
    Actions      :Traffic-rate: 0 bps  (bgp.1)
```

What does the customer see?

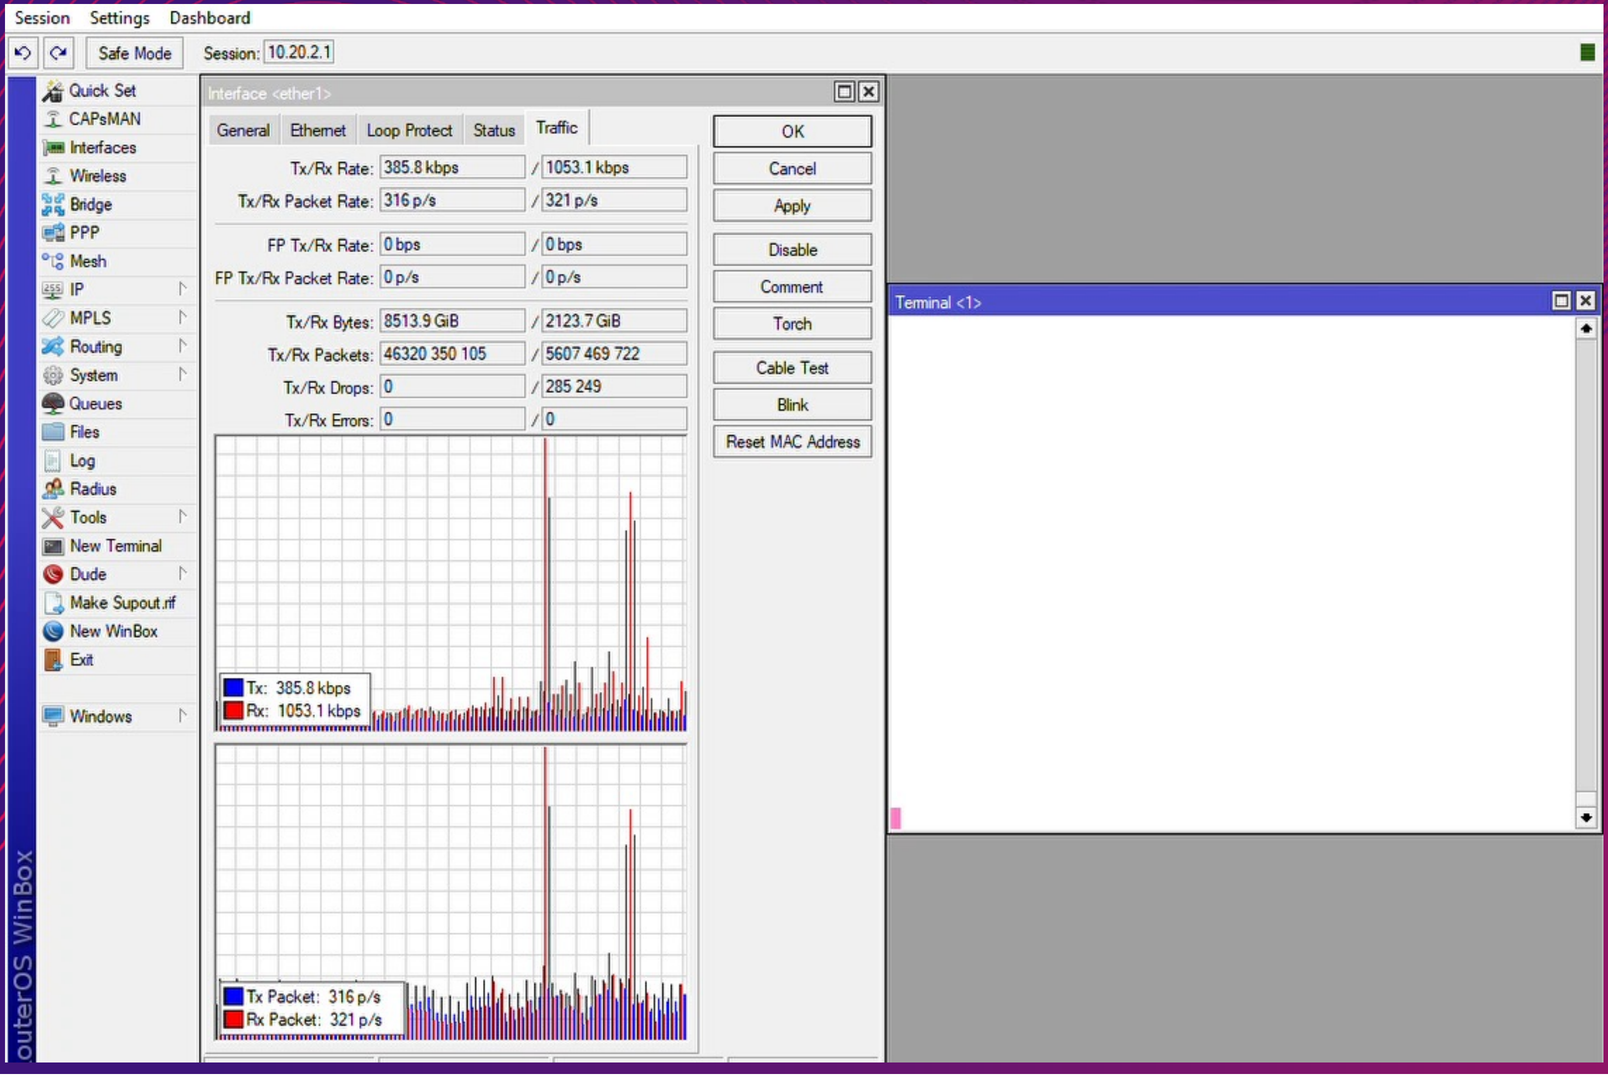

Session: 10.20.2.1

Tx/Rx Rate: 385.8 kbps / 1053.1 kbps

Tx/Rx Packet Rate: 316 p/s / 321 p/s

FP Tx/Rx Rate: 0 bps / 0 bps

FP Tx/Rx Packet Rate: 0 p/s / 0 p/s

Tx/Rx Bytes: 8513.9 GiB / 2123.7 GiB

Tx/Rx Packets: 46320 350 105 / 5607 469 722

Tx/Rx Drops: 0 / 285 249

Tx/Rx Errors: 0 / 0

🟦 Tx: 385.8 kbps
🟥 Rx: 1053.1 kbps

🟦 Tx Packet: 316 p/s
🟥 Rx Packet: 321 p/s

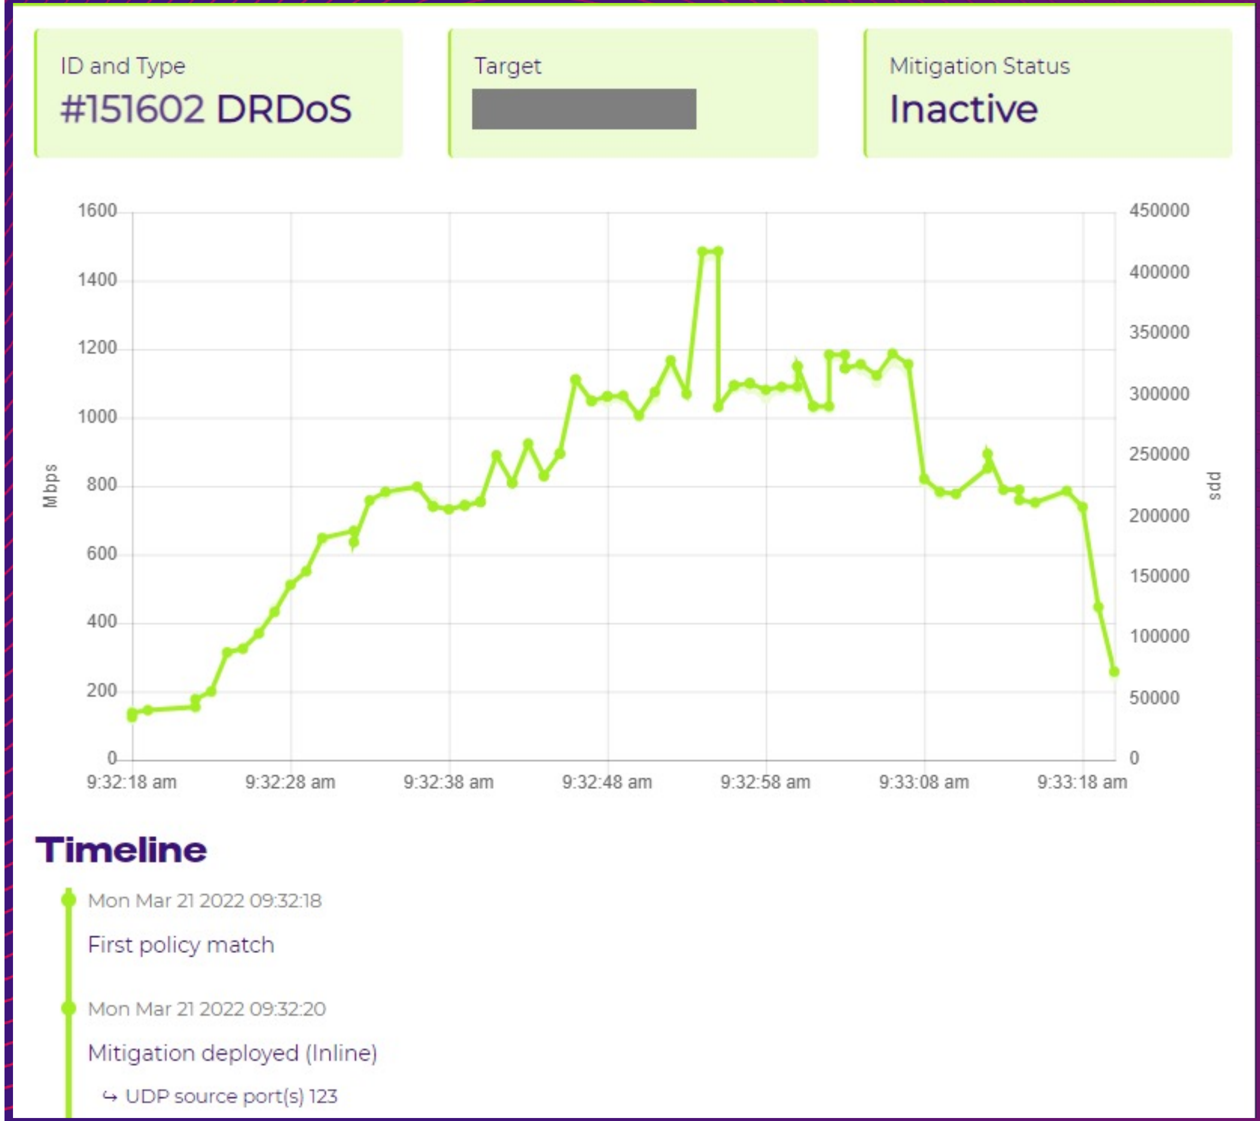

ID and Type

#151602 DRDoS

Target

[redacted]

Mitigation Status

Inactive

Timeline

Mon Mar 21 2022 09:32:18

First policy match

Mon Mar 21 2022 09:32:20

Mitigation deployed (Inline)

↪ UDP source port(s) 123

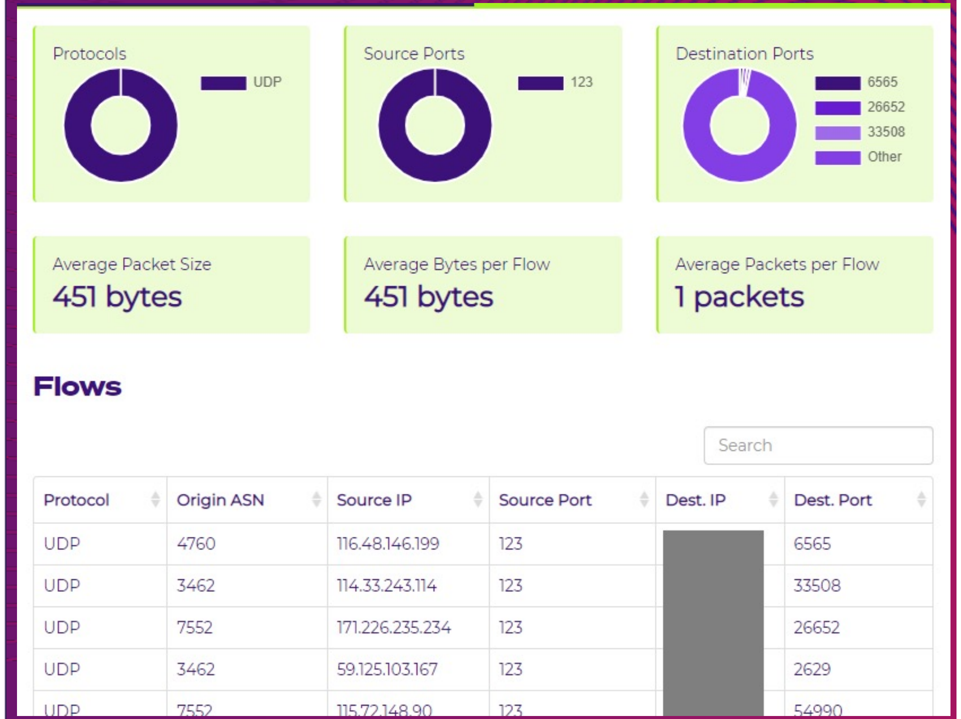

Protocols

Source Ports

Destination Ports

Average Packet Size

451 bytes

Average Bytes per Flow

451 bytes

Average Packets per Flow

1 packets

Flows

Search

Protocol	Origin ASN	Source IP	Source Port	Dest. IP	Dest. Port
UDP	4760	116.48.146.199	123	[redacted]	6565
UDP	3462	114.33.243.114	123	[redacted]	33508
UDP	7552	171.226.235.234	123	[redacted]	26652
UDP	3462	59.125.103.167	123	[redacted]	2629
UDP	7552	115.72.148.90	123	[redacted]	54990

swoop

How does it perform?

We DDoS ourselves once an hour and measure the network "time-to-mitigate"

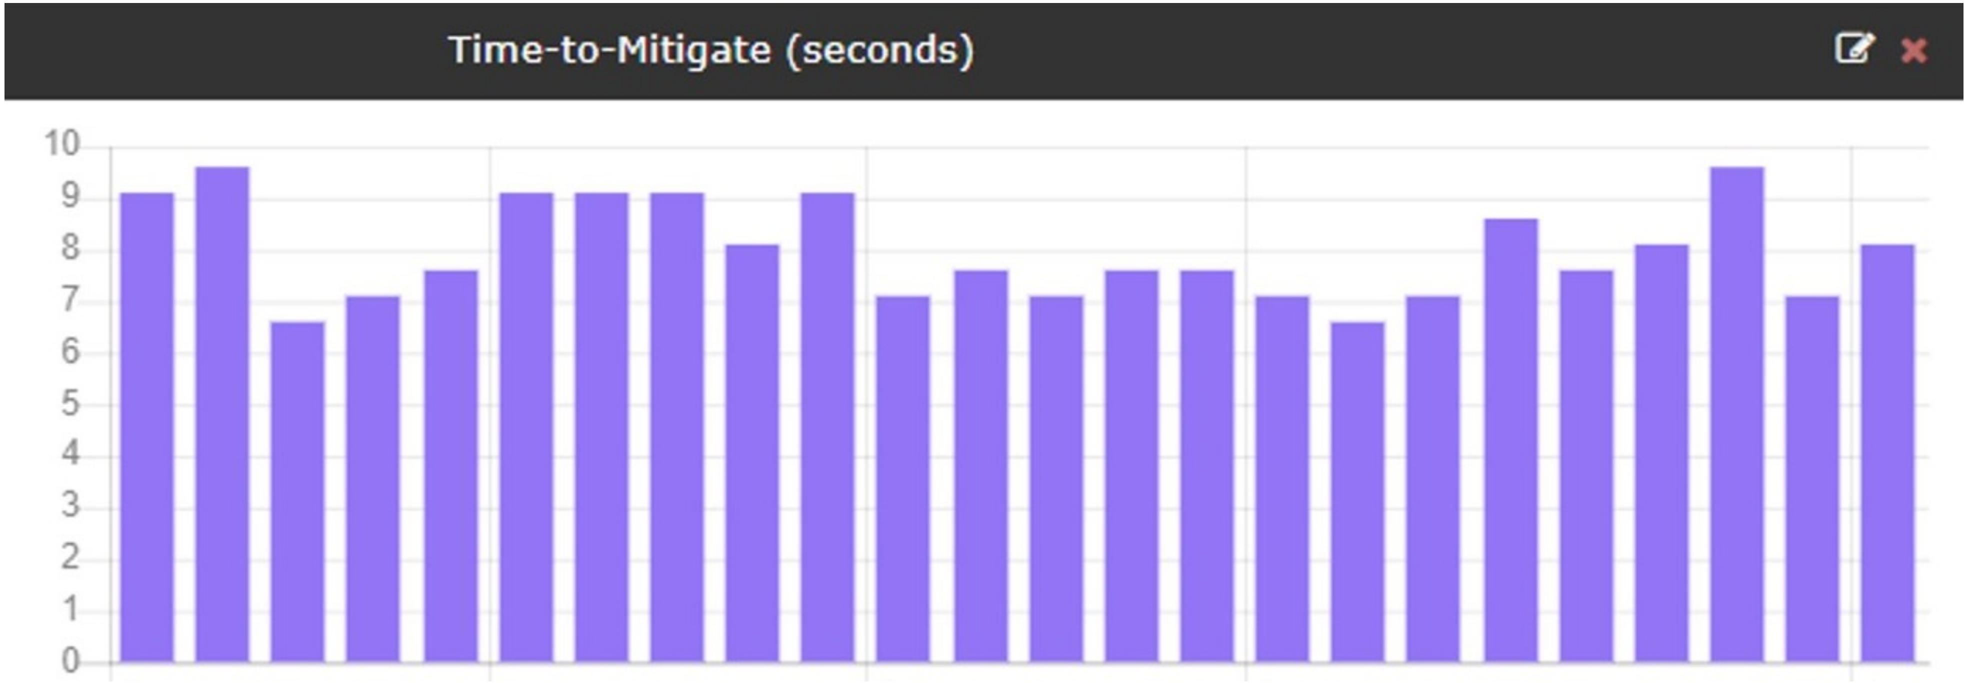

By the numbers

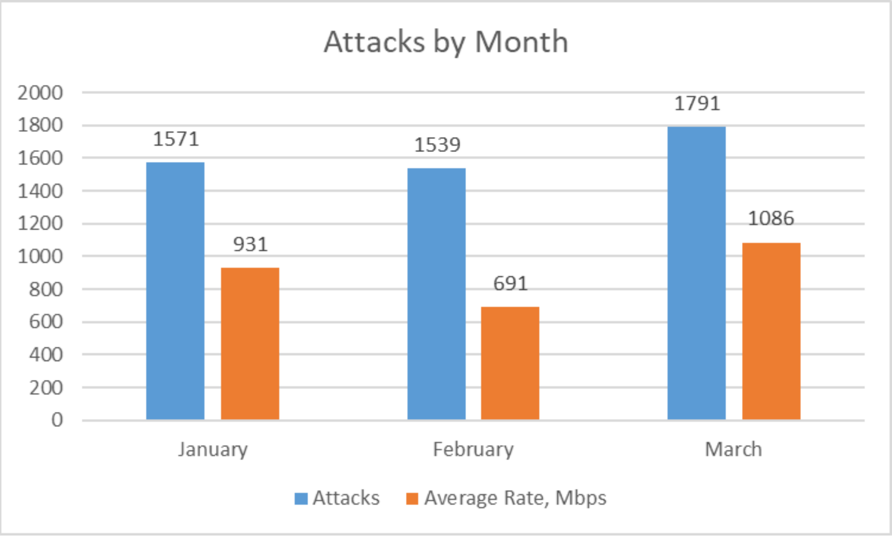

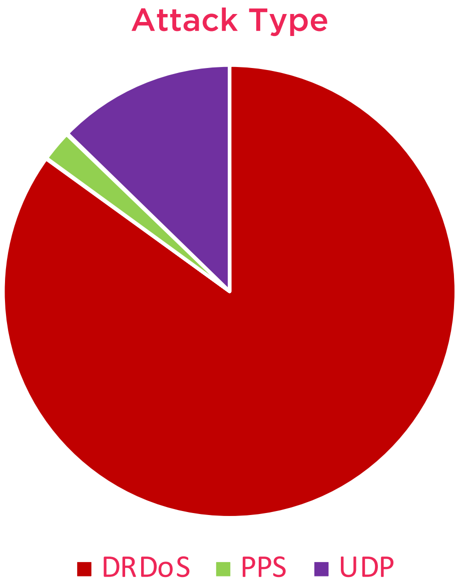

1,600
Avg. attacks per month

900 Mbps
Average attack size

100.2 Gbps
Largest attack to date

Future Improvements

Flow Optimization

July 2022

Policy Customization

December 2022

Better Analytics

December 2022

Swoop 2022

200+ Gbps
Transit Capacity

500+ Gbps
Peering Capacity

50+
Datacentre PoPs

Check out our live network map: swoop.com.au/wholesale/network-map