

Anycast in "The Cloud"

22.10.18
Brett Carr



NOMINET

Agenda

- Introduction
- Short history of our DNS Infrastructure
- Expansion and Cloud choices
- Anycast in the cloud, simple/cost effective
- Problems don't give me no problems
- Futures, where do we go from here



Introduction

Who are Nominet

.uk and GTLD registry operator

RSP for 30+ TLDs

Who am I

Brett Carr, Manager DNS Team

Who was involved

DNS Team:

Karl Dyson

Paul Harris

Alberto Lopez

James Richards

Arife Vural Butcher



Nominet's DNS Infra

- Up to 2015 – 7 unicast nodes
Physical Infrastructure
- 2015 – 8 Anycast nodes UK/EU/US
4 Nameservers
On premise virtual Infrastructure
- 2018 – Expansion ??



Expansion

- Building more nodes is expensive
- Using other peoples computers is cheap(er)
- Cloud Choices
 - AWS
 - Azure
 - Google
 - Others?



Anycast in the Cloud

- AWS selected as the most potentially suitable
- Issues
 - Support for using your own ip space?
 - Load balancers do not support UDP
- Search for help?
- Netactuate/Amazon Direct Connect

Netactuate

- Plenty of experience in Anycast.
- Infra in 25+ Locations globally (more than AWS)
- Solid experience with other DNS providers
- API Access
- Pricing as good as AWS
- Built in DDOS Protection.



VM's in netatcuate

- 4 Locations selected
DFW, GRU, HKG, SYD
- One VM in each location serves all zones
- 8 Vcpu 32gb Memory
- Exabgp (peering with netactuate)
- Health Checker
- dnsdist
- nsd
- Turing collector



Kittens vs Cattle

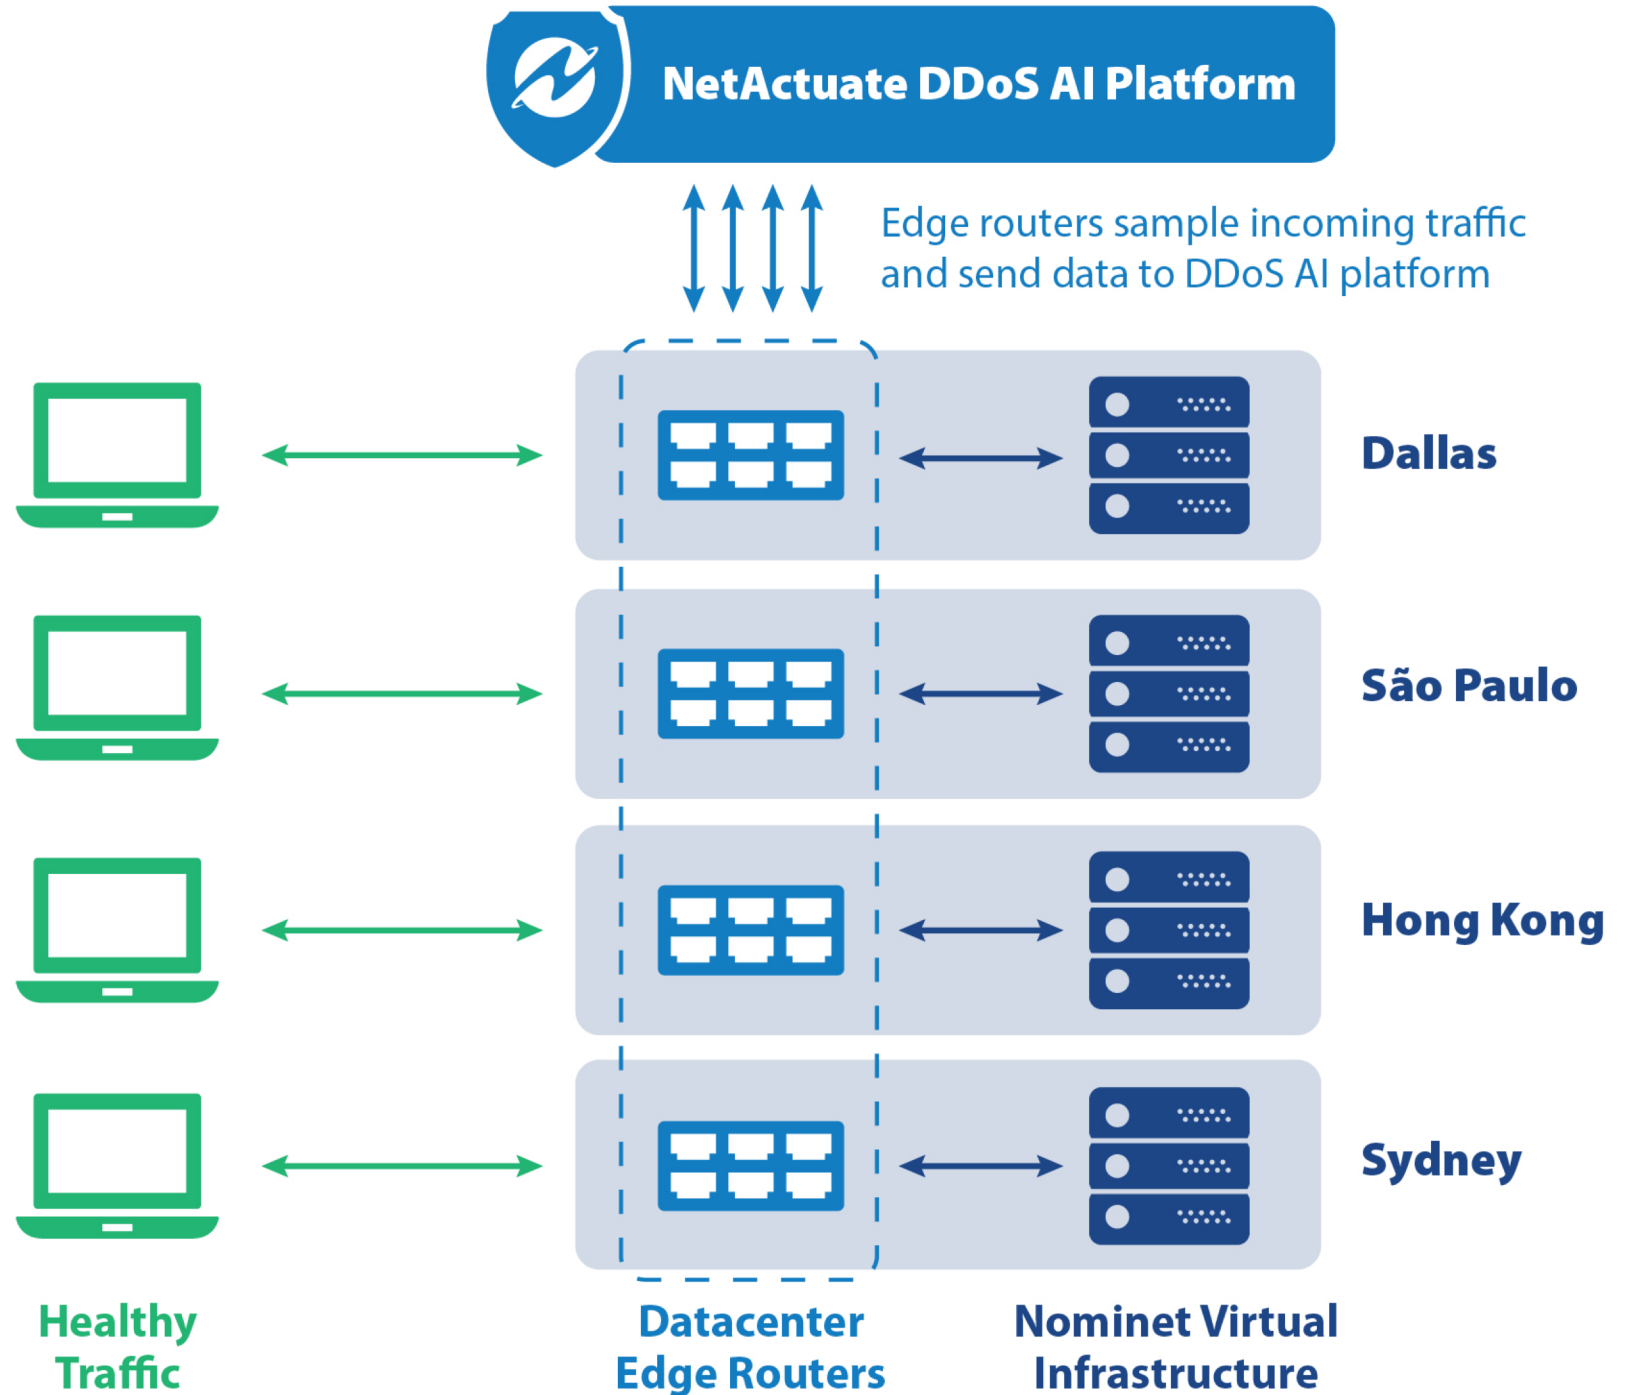
- Immutable Infrastructure
- Built using combination of:
 - Single Image
 - Ansible roles/playbooks
- Only maintain the image
- Birth/Use/Kill/Rebirth



Normal Routing and Monitoring

No Threat Present

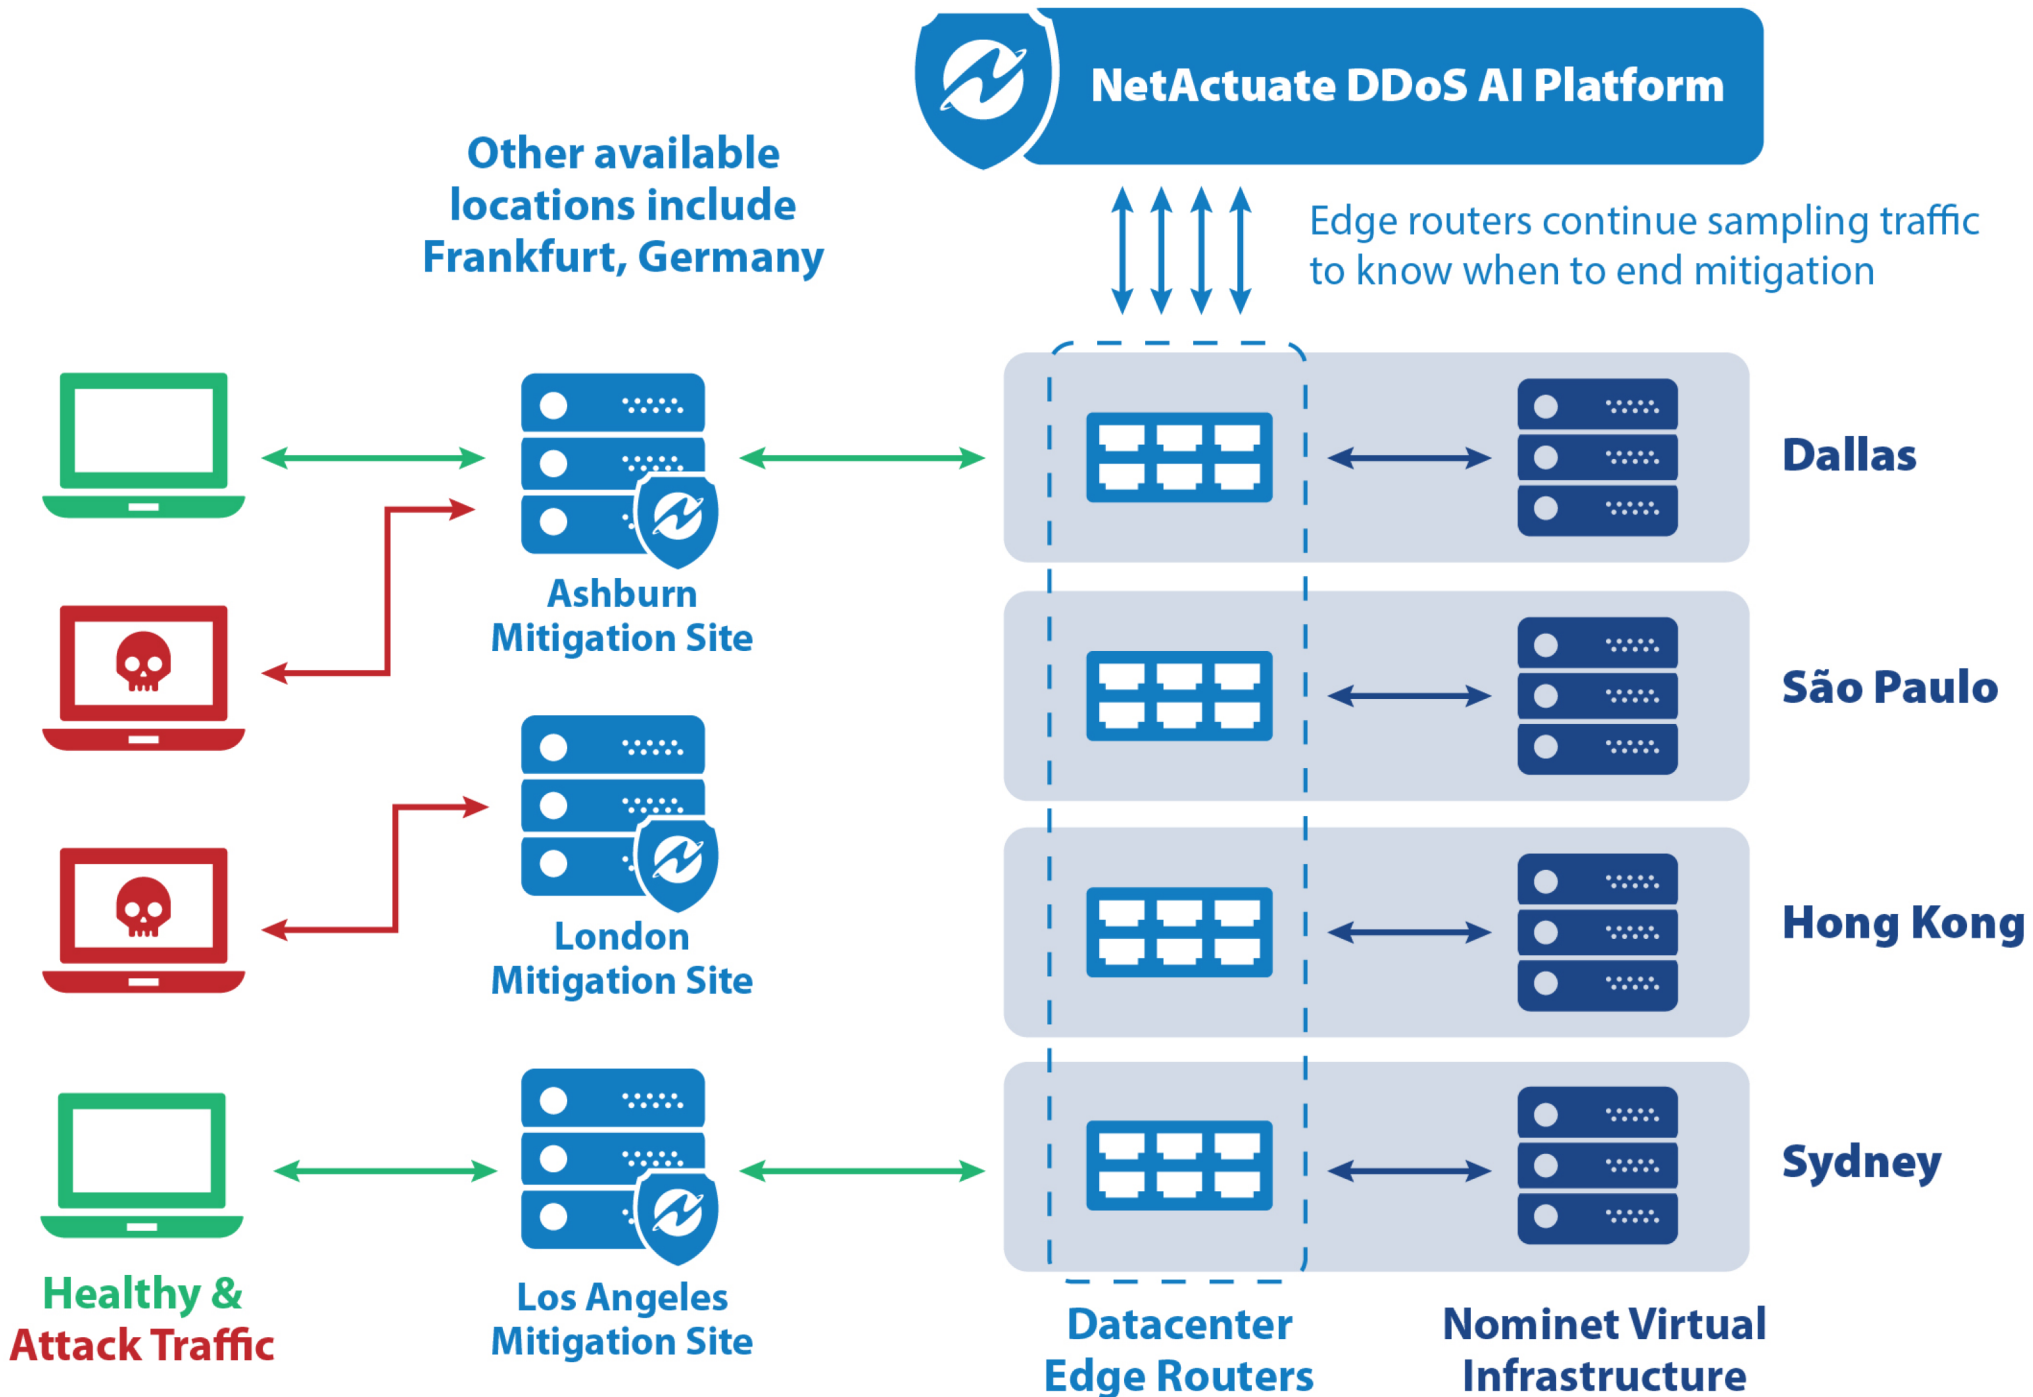
NetActuate's custom-built DDoS AI platform examines all traffic entering Nominet's network. The DDoS AI platform is continually learning Nominet's unique traffic patterns for smarter ongoing protection against attacks.



DDoS Mitigation Active

Threat Detected

When suspicious activity is detected at Nominet's network perimeter, incoming traffic is rerouted to the closest of NetActuate's four global locations for the fastest possible mitigation. These mitigation sites can handle attacks up to 1.2 TB in size.



Problems

- Less traffic than expected
Tweaks made in routing policy by providers
Tweaks made in BGP Config by us
- Global sites can do 500K QPS +
- Cloud sites around 100K QPS
- Care needs to be taken in a DDOS
- TCP Offload needs to be disabled in VM



Futures

- Further roll out into new sites
- Use similar infrastructure for recursive platforms



Thanks for Listening

Questions?

