

Building an ISP in 2017

(using modern methodologies)

andrew@aussie.net

AusNOG 2017 - 7st September 2016

Who is Andrew Khoo?

- I have been providing Internet access since 1992
- I started seeding exchange points in east-coast Australia since 1997, around the same time as WAIX
- Teaching BGP to ISPs since 1998
- Global Network Architect - Akamai early 2000s
- Architecting OTT/media solutions for a large telco last few years
- Building a greenfields RSP in 2016/7...

Legal Disclaimer:

Examples provided in this presentation may or may not relate to the work I am doing for amaysim.

Nothing discussed or presented should be interpreted as a forward-looking statement.

All opinions are purely personal.



What is the big picture?

A modern green-fields ISP in 2017 does not resemble one in 2007, or 1997.

Or even 2012.

NBN is here to stay - warts and all.

- CVC has shifted focus from last-mile engineering, to engineering for the last-mile.

Consolidation is happening

- In 30/06/17 Quarterly Wholesale Market Indicators report, only Aussie Broadband is mentioned outside the “big 4” for all access types except satellite

The opportunity?

- Mobile MVNO with large customer base for cross-selling
- Large pool of talented developers and data analysts
- NBN creating “forced churn” event

The risks?

- Marketing against “big 4” and their downstream resellers
- Build vs Resell – the risk/reward challenge

Design challenge

- Use modern methodologies to implement systems
- IPoE for customer termination
- IP only - MPLS is the ATM of this century
- Certify for NBN inter-operability in 1Q
- Roll out national footprint in 2Q

A paradigm shift in hardware


- Terminate IPoE sessions on a modern L2/L3 switch
 - Untag double-tagged NBN sessions to L3
 - Run dhcp server on switch
- Carry all IX and direct peering prefixes in same L2/L3 switch
 - Able to process multiple full views. Quickly.
- Build to 100GbE “core” – 40GbE/10GbE for NNIs
 - 1Gbps is dead. Long live 1Gbps.
- Small footprint, low power consumption stack
 - It’s not about the size of the kit
 - Power is \$\$\$

The hardware candidates:

- Cisco ASR9k
- Cumulus on Edgecore/Accton
- Arista 7050/7280

A (legacy) spanner in the works: PPPoE!

A small stack is a good stack

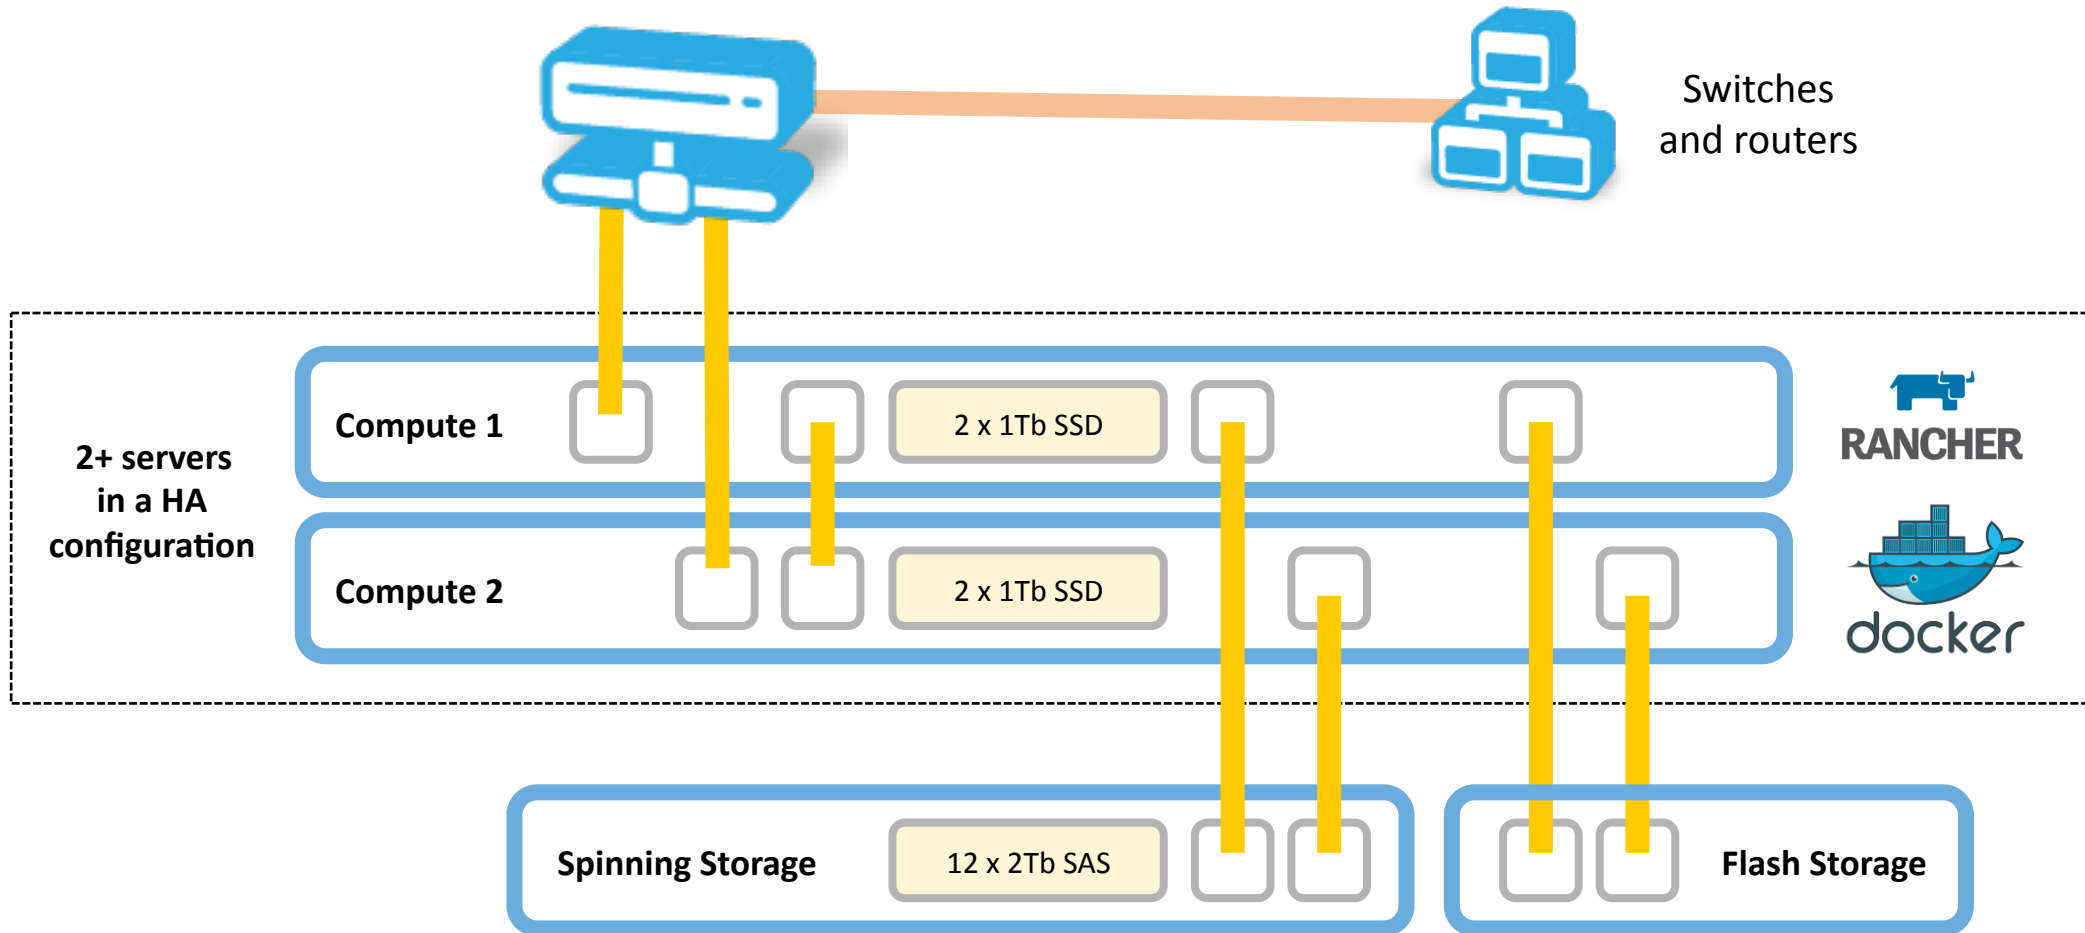


...Aristas with a sprinkling of Mikrotik

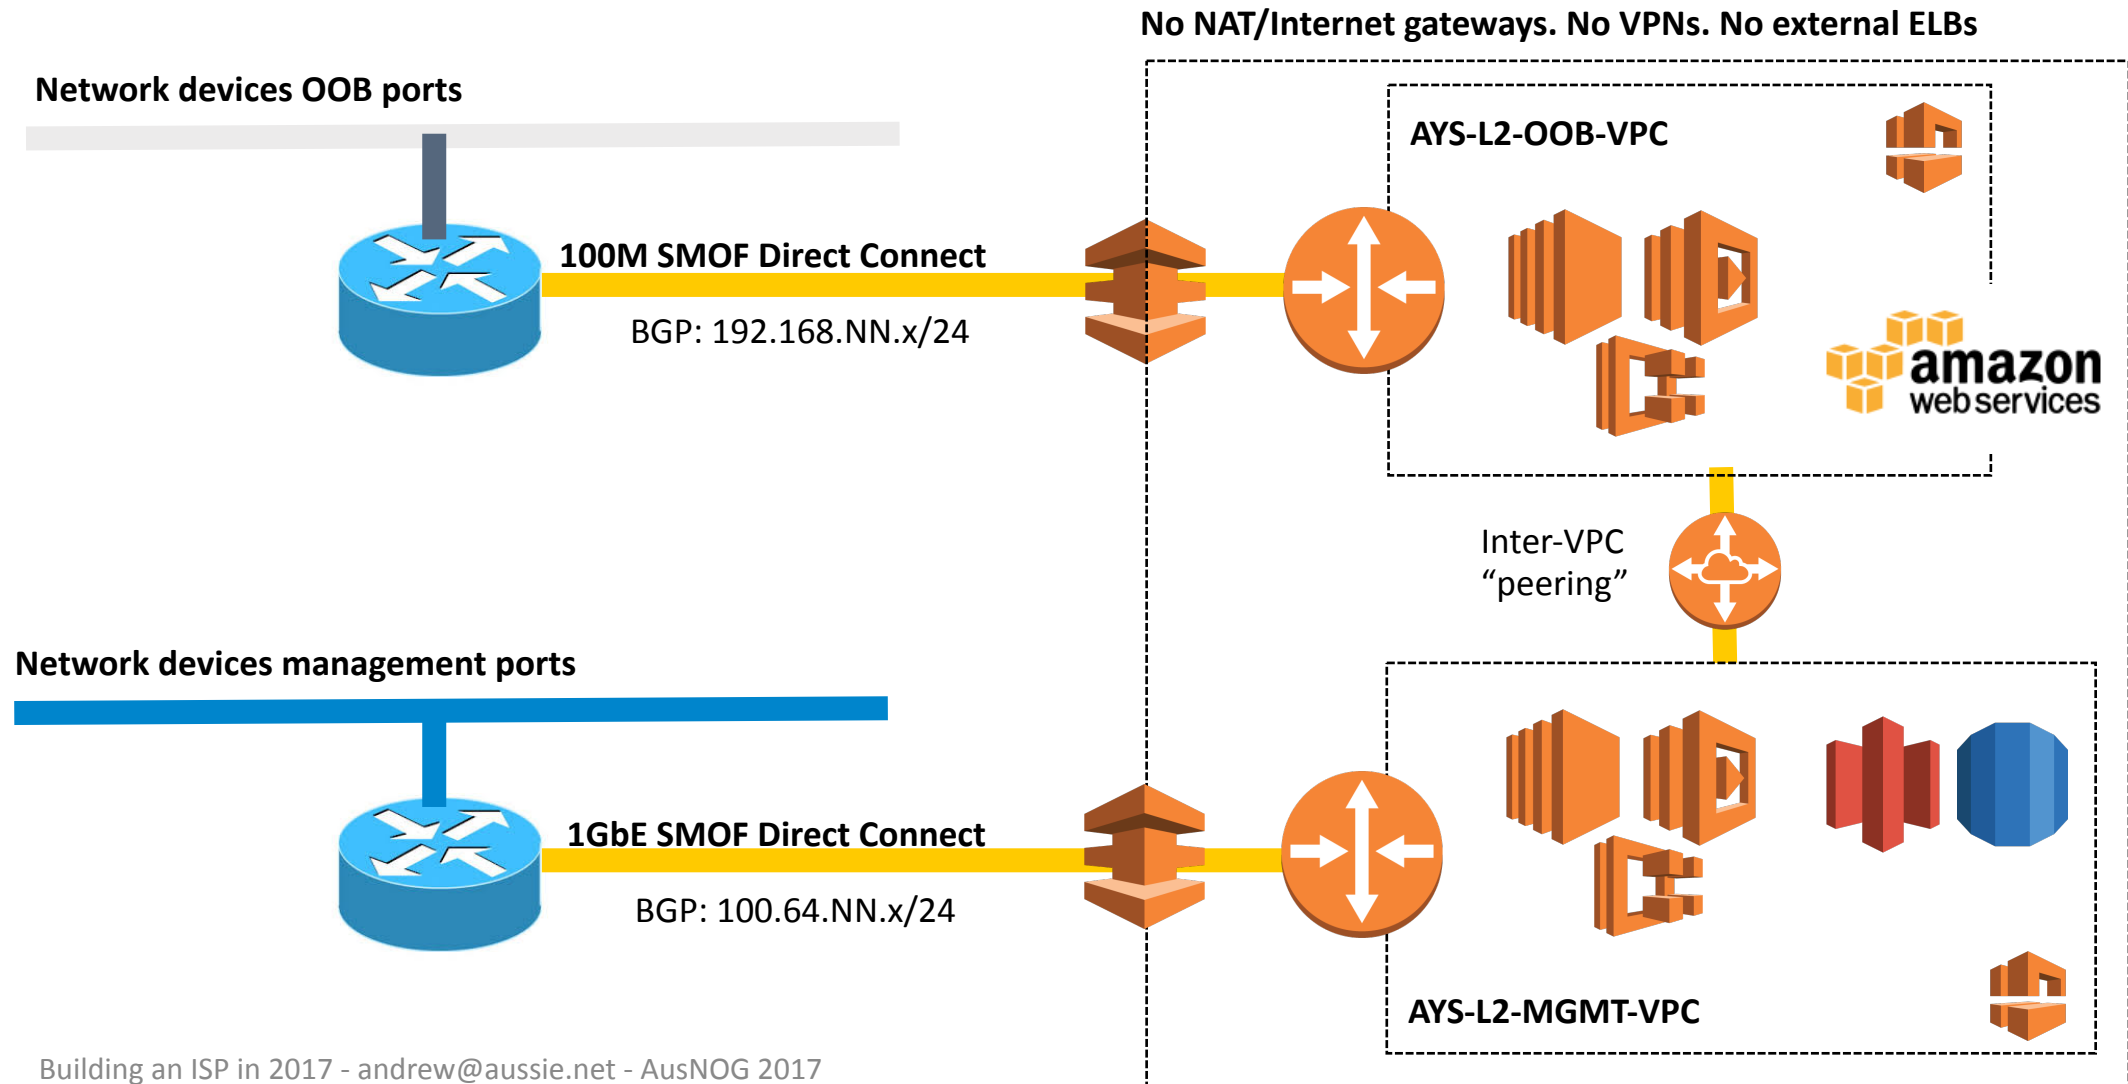
A paradigm shift in software

- Code as infrastructure
 - serverless is ideal and desired state
 - containers for long-held sessions
- Everything – absolutely everything – must use API glue
 - Modular design = parallel development
 - Avoids religious battles over language choice
- IPv6 first
 - Regard legacy Internet as IPv4-as-a-service
- Leverage “cloud” when it makes sense
 - Running everything in AWS/GCP/Azure/Aliyun will make you go broke

POP Compute Cluster



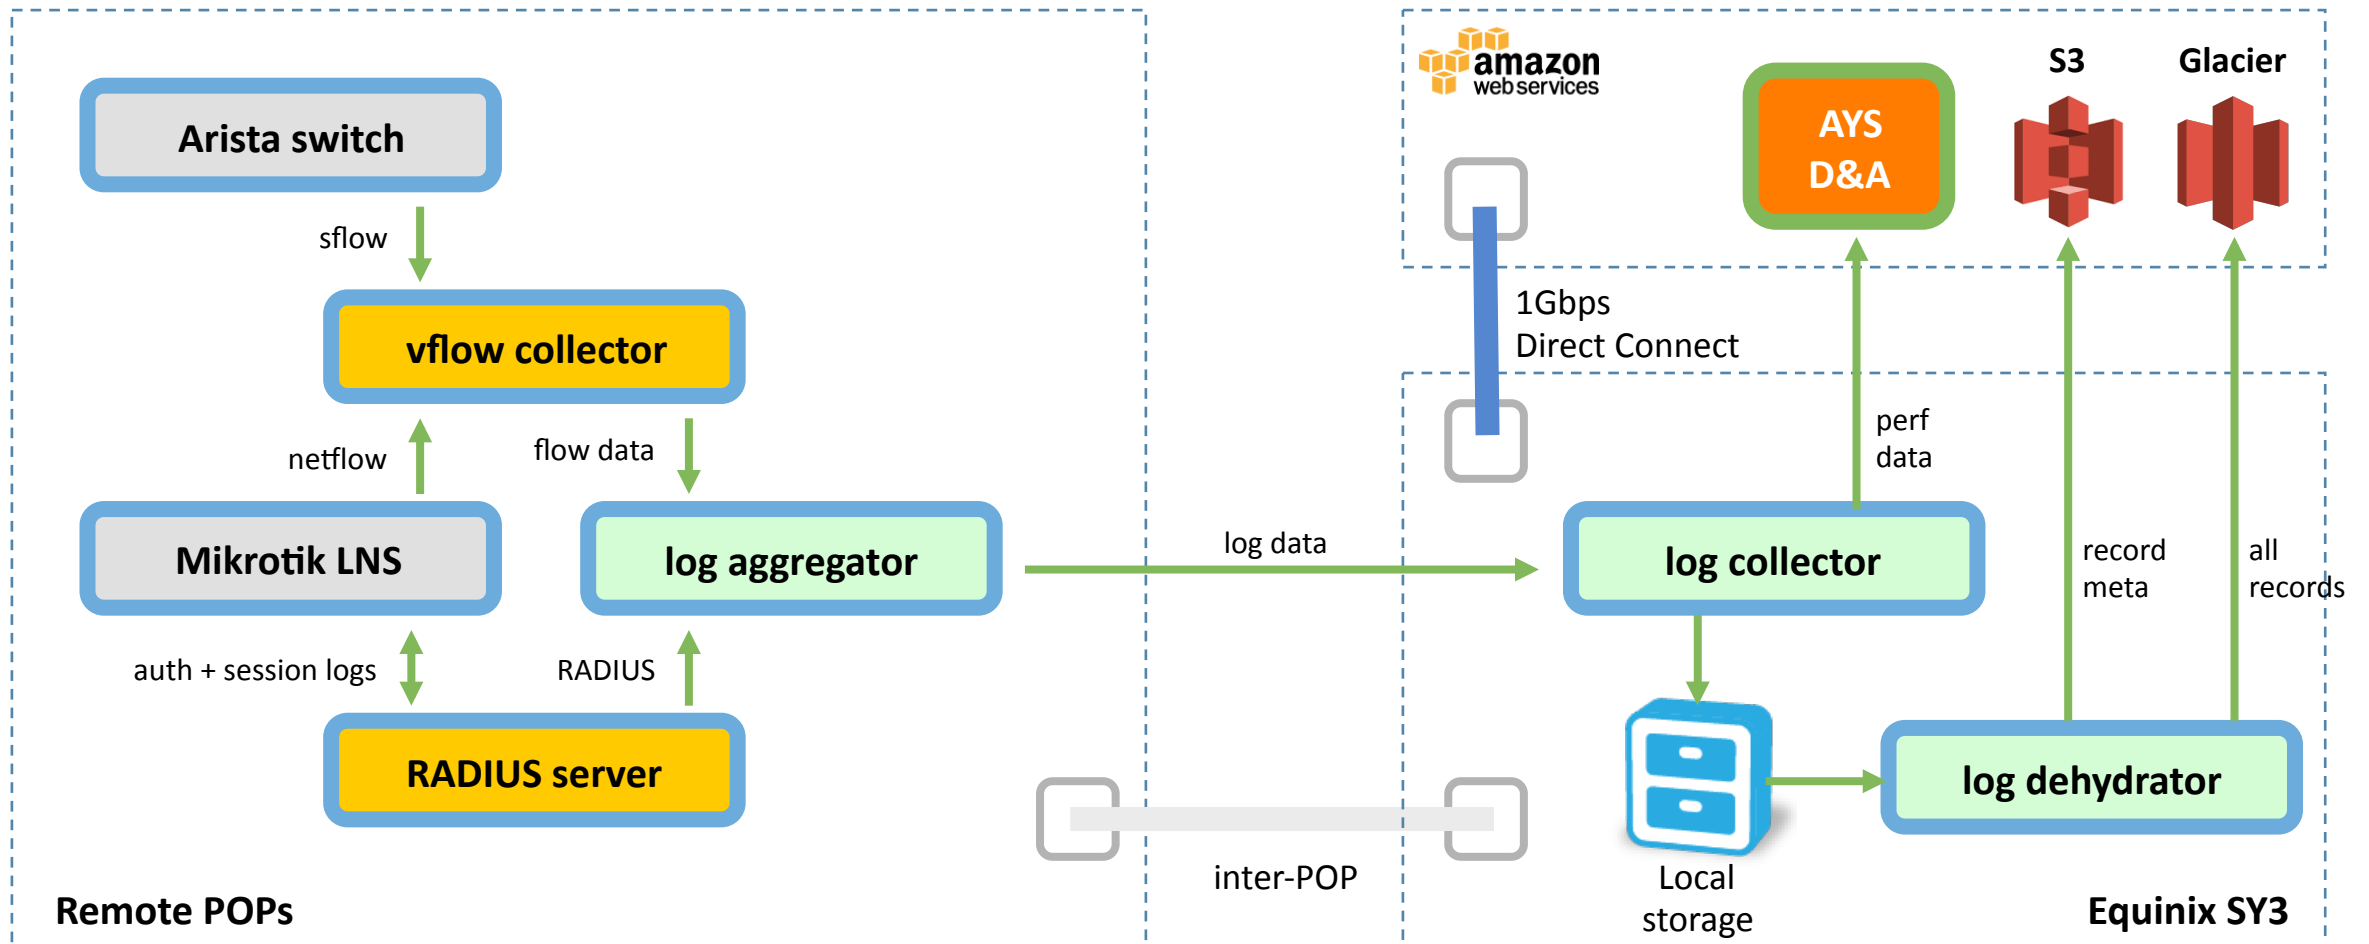
Use “cloud” where it makes sense

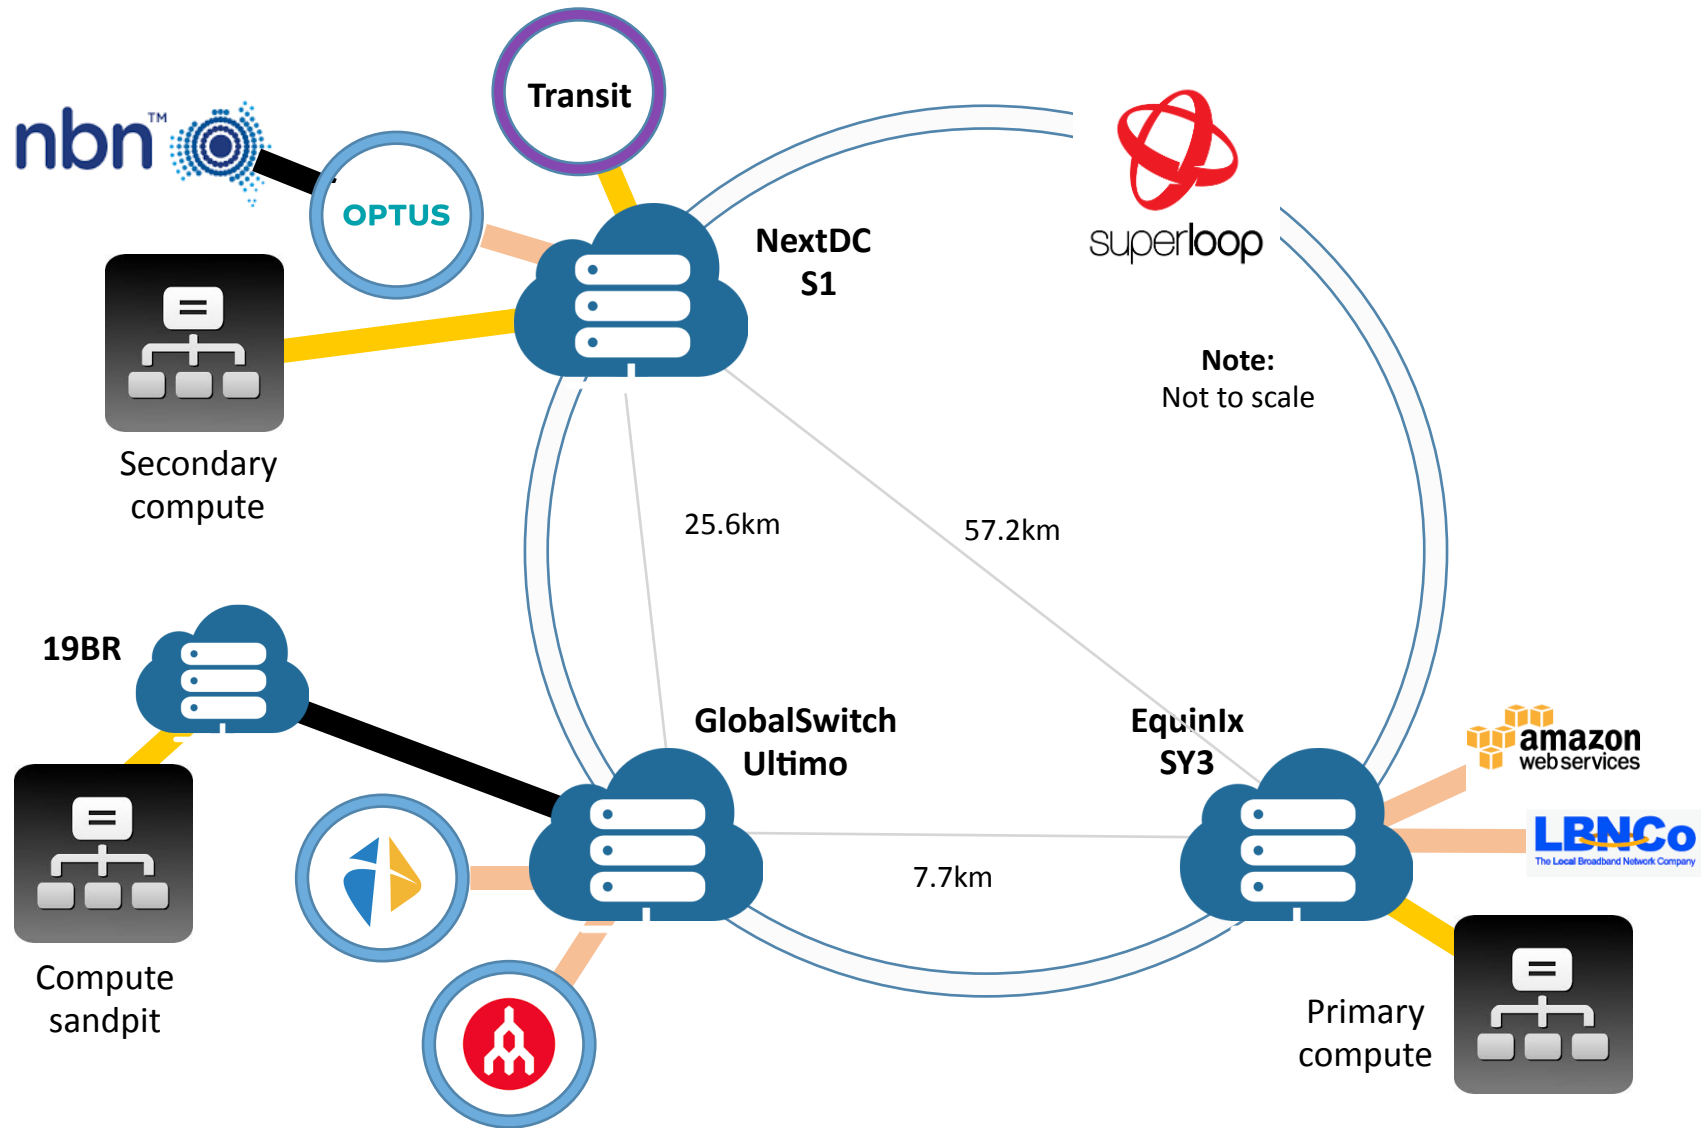


Scary stuff for network engineers

- Everything runs stateless
 - Pull configs on start
 - Rollback-on-failure
 - No SSH access in production
- Closed feedback loop for auto-configuration
 - Telemetry is king – automate decisions
- Engineer for lossless logs and telemetry (flow data)
- All the majors are already doing this – SREs etc

Data Retention and QoS Workflow





Our full tech stack is consistent...

Modern software development

- Docker/CoreOS as container host
- GoCD for continuous deployment
- Rancher for orchestration
- Lambda for serverless
- Terraform for build orchestration
- Packer to create AMIs
- Consul for service discovery

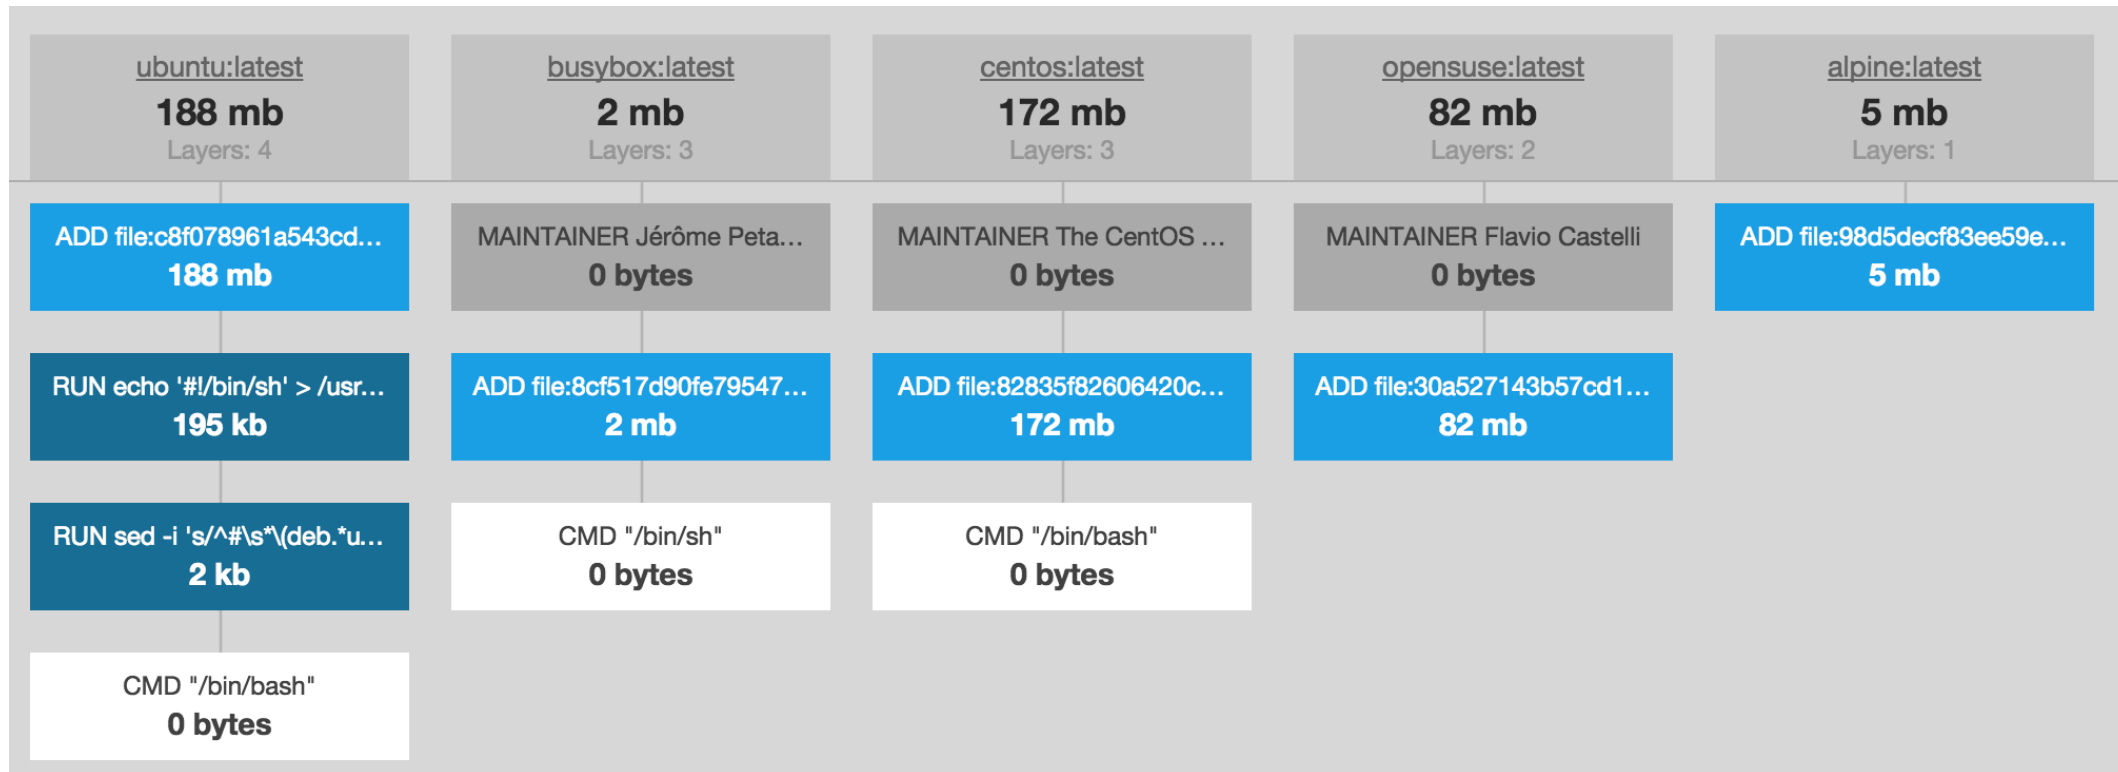
Moving hardware world to unit testing

Our engineering principles:

- Deliver, continuously
- Integrate, continuously
- Improve, continuously
- Automate all the Things
- Monitor all the Things



We instantiate containers in 5mb..



Small is good when running code on a switch

Tooling used for “hardware” automation

- SaltStack – client/server automation & orchestration
- Ansible – provisioning via SSH
- GitHub – repo for all configs that are pulled
- Jenkins – uses GitHub webhook to trigger
- Kong – API gateway (auth + logging etc)
- RabbitMQ – the ruler of the entire stack

Examples of automated deploys

- update radiusd and/or dhcpd with new auth
 - pull new AVCID information or PPPoE auth creds to remote stacks
- provision a new end-point
 - set up port on switch – VLAN, VXLAN, br0
 - activate virtual port and configure link-local
- bring up new BGP peer
 - verify prefixes against PeeringDB to release from “quarantine”
 - create route-maps, prefix-lists and apply communities
- add/remove BYO customer IPv6 prefix to ACLs
 - triggered via user portal
 - add prefix to ACLs and optionally firewall rules
 - update Radius and DHCPv6 config

Q&A

andrew@aussie.net

