Management Plane in Distributed Infrastructures
A Historical View

Marcio Saito

Sep 2019 – Santa Clara, CA

Once upon a time…


S 7.6 0.2 1448:00 Postman
S 6.9 0.5 1263:36 Postman
S 4.3 0.2 869:32.83 Postman
S 3.6 0.8 877:05.63 Xorg
S 3.3 0.9 145:25.68 chrome
S 3.3 4.4 1718:11 gnome-shell
S 2.6 1.6 136:42.13 chrome
S 2.6 0.0 677:25.06 irq/133-nvidia
S 2.0 0.5 24:16.11 chrome
S 1.3 1.0 17:09.57 chrome
S 1.3 0.8 30:32.55 Web Content
S 1.3 0.1 48:41.34 pulseaudio
S 1.3 0.5 17:28.36 chrome
S 1.3 0.5 26:45.50 chrome
S 1.0 0.0 244:12.57 nvidia-modeset
S 1.0 0.0 0:01.09 top
S 0.5 24:59.28 chrome
22:58.84 chrome

16-port Terminal Server (1995)

YAHOO!®

opengear

ONF CONNECT


"Crash Cart"

Out-of-Band Management

"

An **Out-of-Band Management System** provides a secure alternate path so that a **network engineer** can reach the **console port** of any **network device** even when the production network is disrupted.

Virtualization and Hyperscale Data Centers

Collapsed Data/Management planes

Google Cloud Networking Incident #19009

The network congestion issue in eastern USA, affecting Google Cloud, G Suite, and YouTube has

Incident began at **2019-06-02 11:45** and ended at **2019-06-02 15:40** (all times are **US/Pacific**).

Go… e be… plane be… - aft… was withdrawn, resulting in the significant reduction in network capacity observed by our services and users, and the inaccessibility of some Google Cloud regions. End-user impact began to be seen in the period 11:47-11:49 US/Pacific.

Google engineers were alerted to the failure two minutes after it began, and rapidly engaged the incident management protocols used for the most significant of production incidents. Debugging the problem was significantly hampered by failure of tools competing over use of the now-congested network. The defense in depth philosophy means we have robust backup plans for handling failure of such tools, but use of these backup plans (including engineers travelling to secure facilities designed to withstand the most catastrophic failures, and a reduction in priority of less critical network traffic classes to reduce congestion) added to the time spent debugging. Furthermore, the scope and scale of the outage, and collateral damage to tooling as a result of network congestion, made it initially difficult to precisely identify impact and communicate accurately with customers.

As of 13:01 US/Pacific, the incident had been root-caused, and engineers halted the automation software responsible for the maintenance event. We then set about re-enabling the network control plane and its supporting infrastructure. Additional

https://status.cloud.google.com/incident/cloud-networking/19009

Now, let’s look into the future…

Out-of-Band Management

"

An **Out-of-Band Management System** provides a secure alternate path so that a **network engineer** can reach the **console port** of any **network device** even when the production network is disrupted.

Resilience Fabric

"

A **Network Resilience Fabric** extends the reach of other **Management Systems** at a central location to access **Remote Infrastructure** even if the production network is disrupted or not yet established.

CORD®

Mission

To Build an Open Reference Implementation of CORD®

Our mission is to bring datacenter economies and cloud agility to service providers for their residential, enterprise, and mobile customers using an open reference implementation of CORD with an active participation of the community. The reference implementation of CORD will be built from commodity servers, white-box switches, disaggregated access technologies (e.g., vOLT, vBBU, vDOCSIS), and open source software (e.g., OpenStack, ONOS, XOS).

Network connectivity: user / data plane

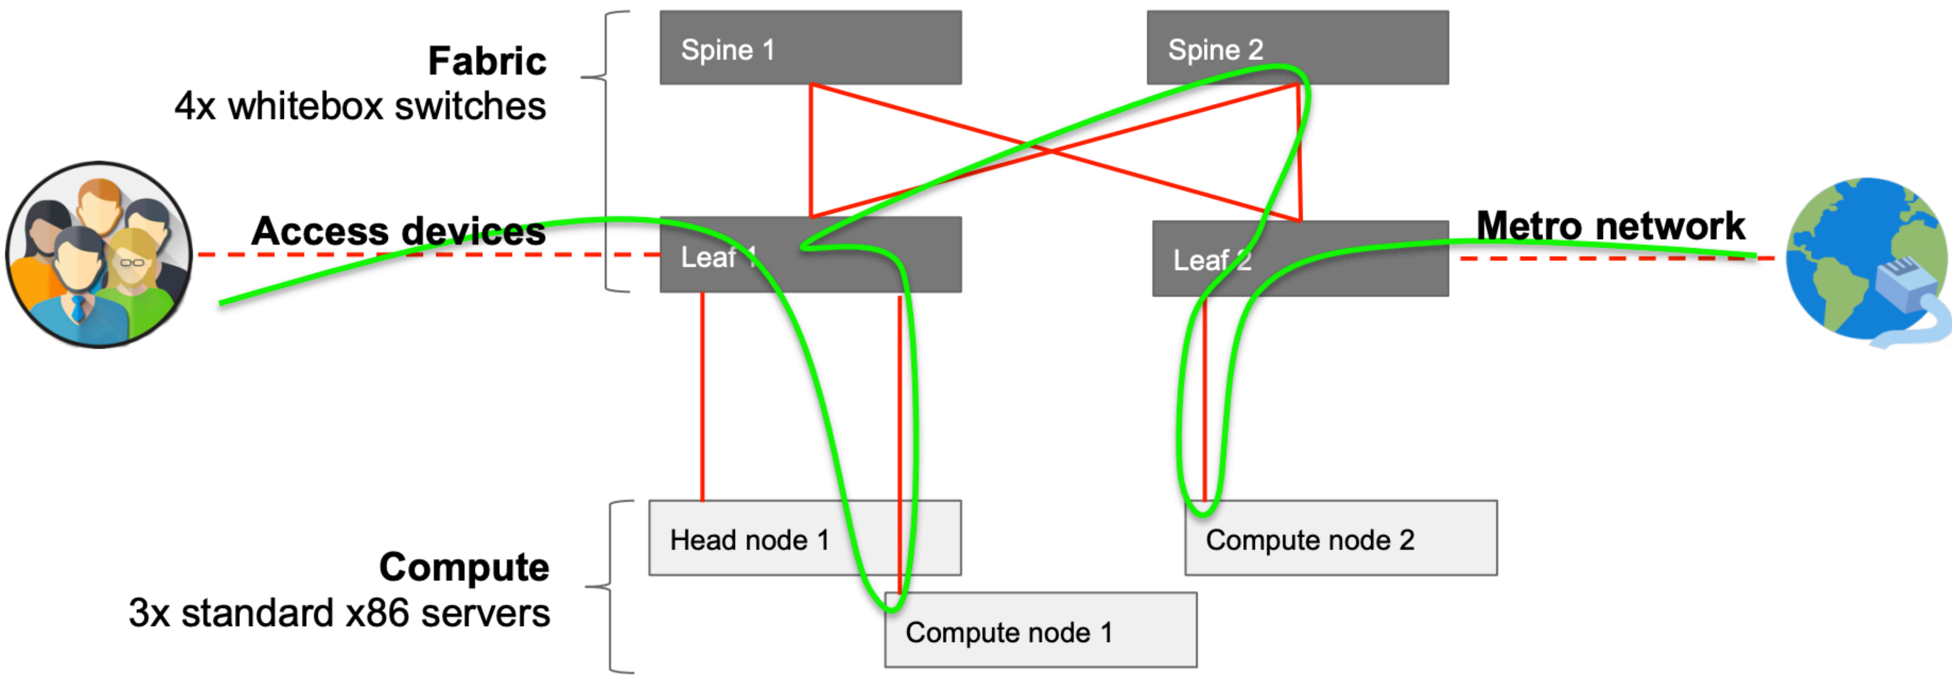

#OpenCORD

Network connectivity: a complete view

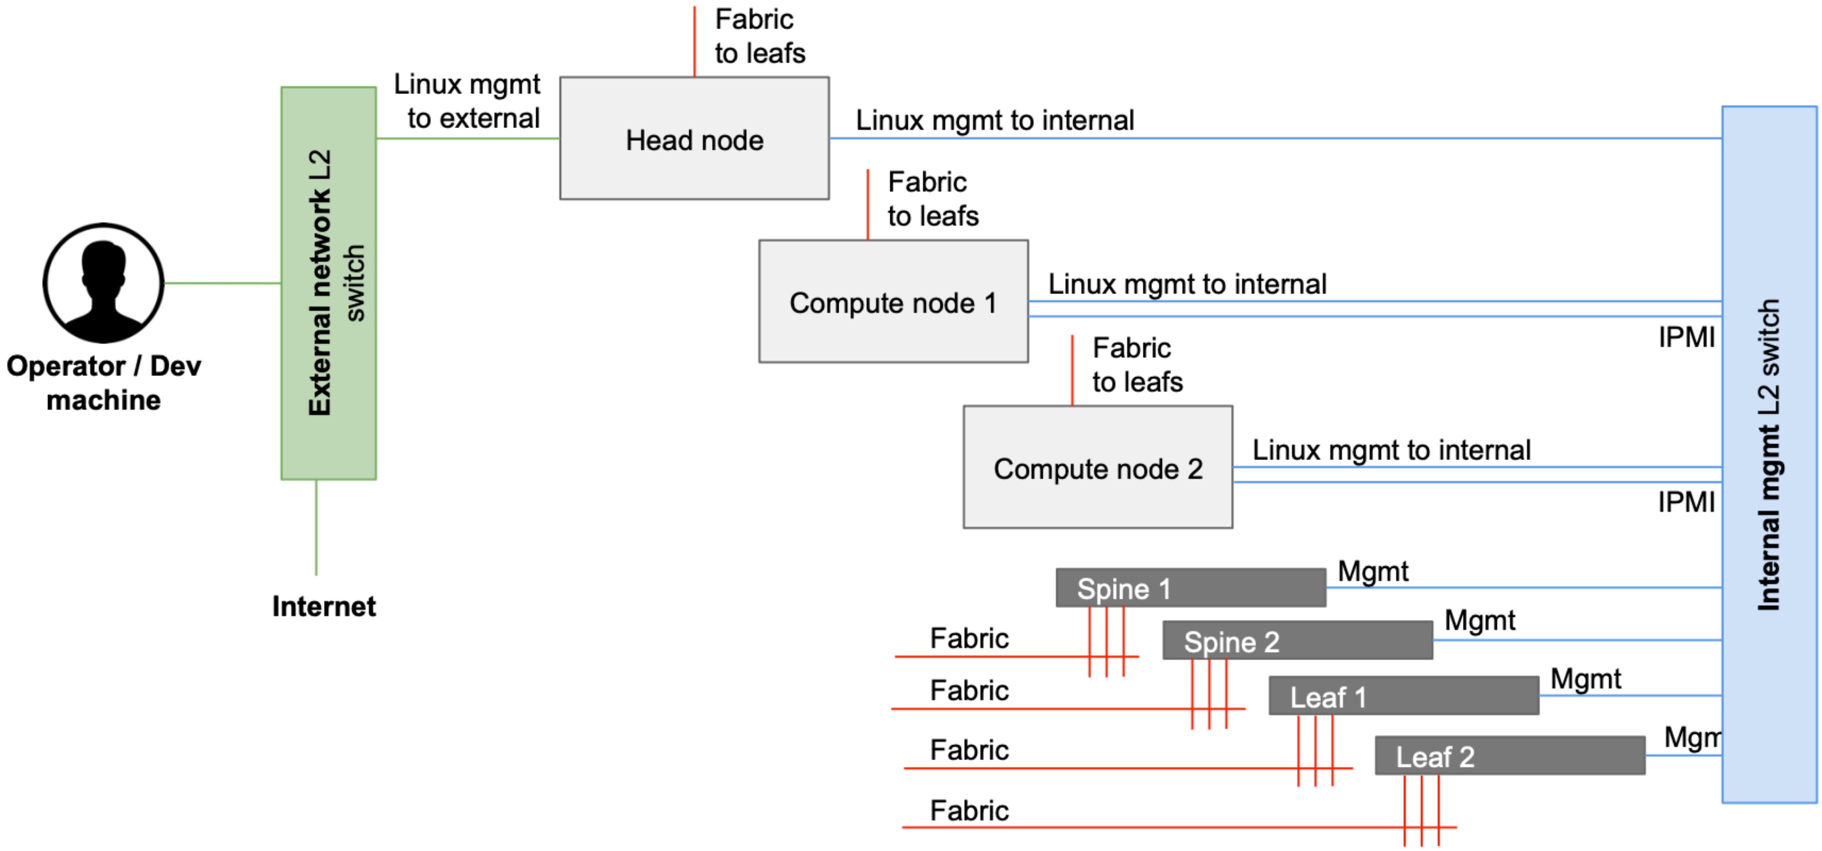

To Architects of CORD

Management beyond connectivity and Day-One Provisioning

- As network management becomes automated and distributed, there will be a need to deploy management applications at the edge for provisioning, monitoring, configuration management, telemetry
- Consider extending current definition of Head Node (currently an access gateway and provisioning server) to also include the ability to host post-provisioning management applications (or explicitly define compute nodes dedicated to run management applications)
- It is no longer about connecting humans to ports. It is about extending the reach of management systems (preferably transparently) to the edge of the network. Consider defining an “automation gateway” that standardizes how that is done

To Implementers of CORD
Keep management plane separation in mind

- The collapsing of management and data planes into the same infrastructure drastically reduces the availability, resilience and security of edge infrastructure
- Implementation-time decisions can affect the separation of planes even when the basic architecture design is solid
- When implementing CORD, ensure the separation is not only logical, but includes the physical infrastructure, network links, and management applications to ensure resilience from the edge to the core

Thank You

Marcio Saito

Marcio.Saito@Opengear.com