Telematics Container Cloud Platform

TCCP Architecture & Deployment

2018-03-05: Rick Davis

verizon connect

The power of Verizon

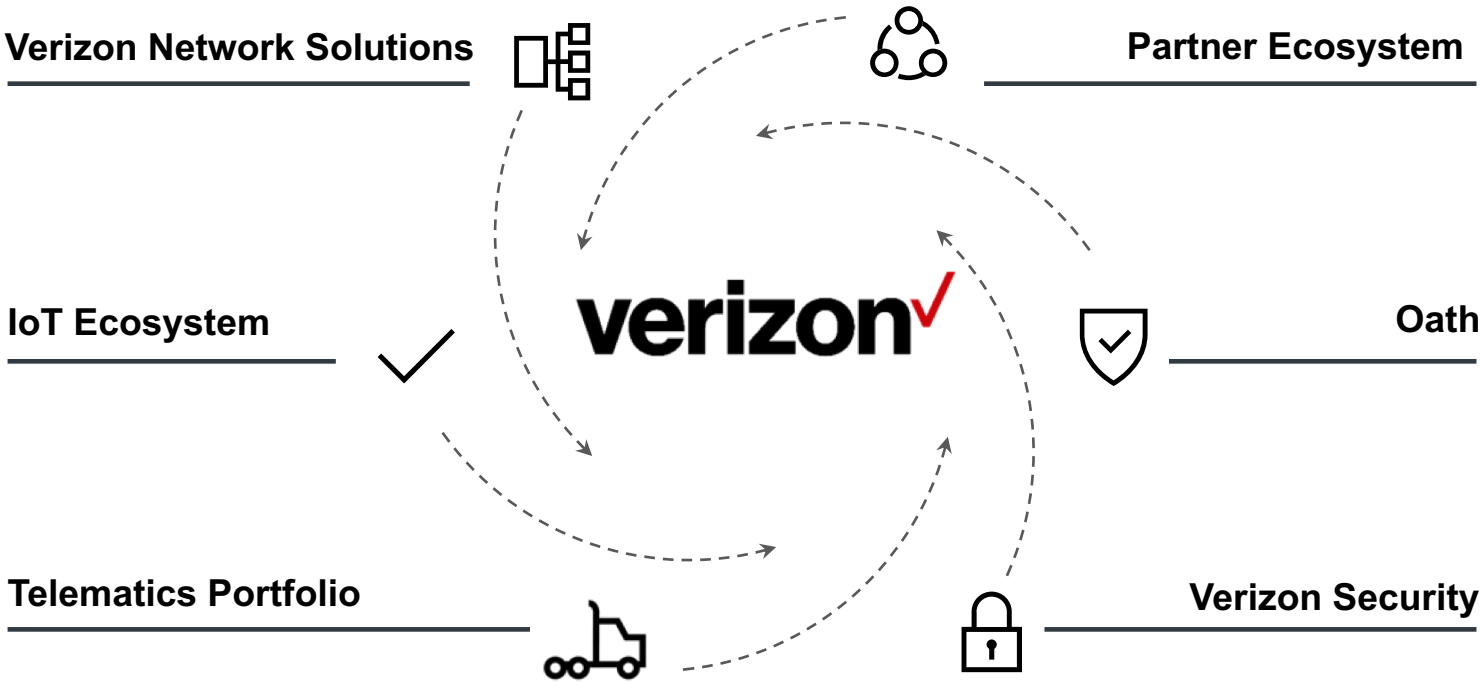

verizon connect

Our technology portfolio is second to none

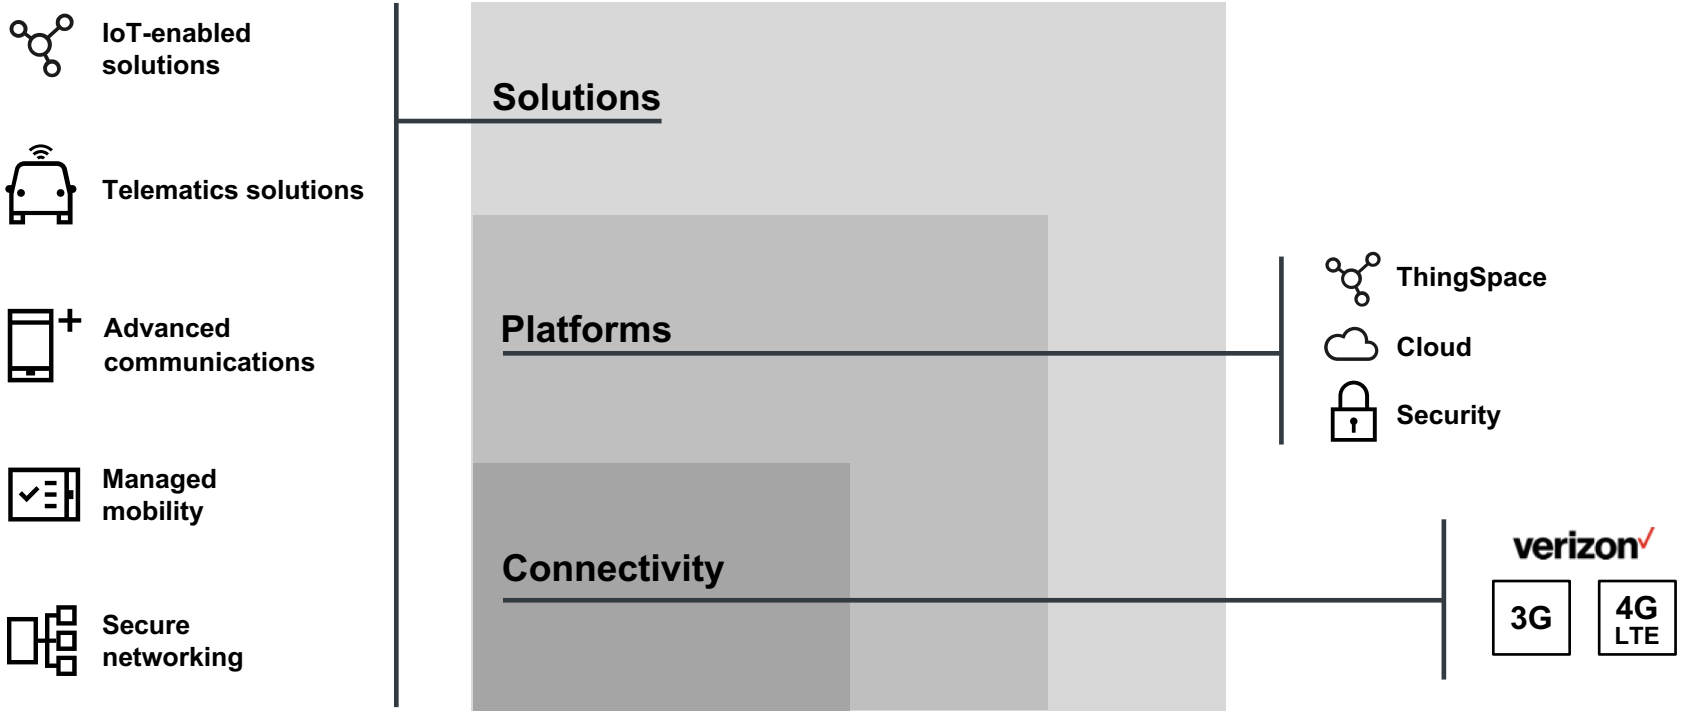

verizon connect

Agenda

1. Goals
2. Architecture
3. Automation
4. Collaboration & Lessons Learned
5. Q&A

verizon connect

Goals

What do we want to do?

verizon connect

Business Drivers

Increase Development Velocity

- Monolith to Microservice
- Latest Software Frameworks
- Containerized
- 100% CICD

Increase Infrastructure Velocity

- Site Rebuild & App Deploys < 12hr
- Repeatable/Predictable
- 100% CICD

Future Proofing

- Flexible & Extensible
- Easily Scalable
- Cloud Adaptable & Agnostic

verizon connect

Infrastructure Goals

Architecture

- Cost Conscious
  - Open Hardware
  - Open Source Software, Commercial Support Optional
- Global Deployment Capability w/Vendor Support
- Public & Private Cloud Adaptable

Automation

- Server & Network Provisioning
- Integrated Validations
- Infrastructure-as-Code

Extensibility

- Platform with Swappable Components
- Avoid Lock-in, Everywhere
  - Hardware, Software, and Cloud

verizon connect

Architecture

How are we going to do it?

verizon connect

Architecture

eBGP Clos Fabric

- One Protocol/Simplified OPS
- Quagga with ECMP on Hosts
- Anycast Support

Spine/Leaf Topology

- Dedicated Storage Network
- Dedicated Border/Peering

Simple Scalability

- 100's of Racks/Site

Open Hardware & Software, Commercial Support Optional

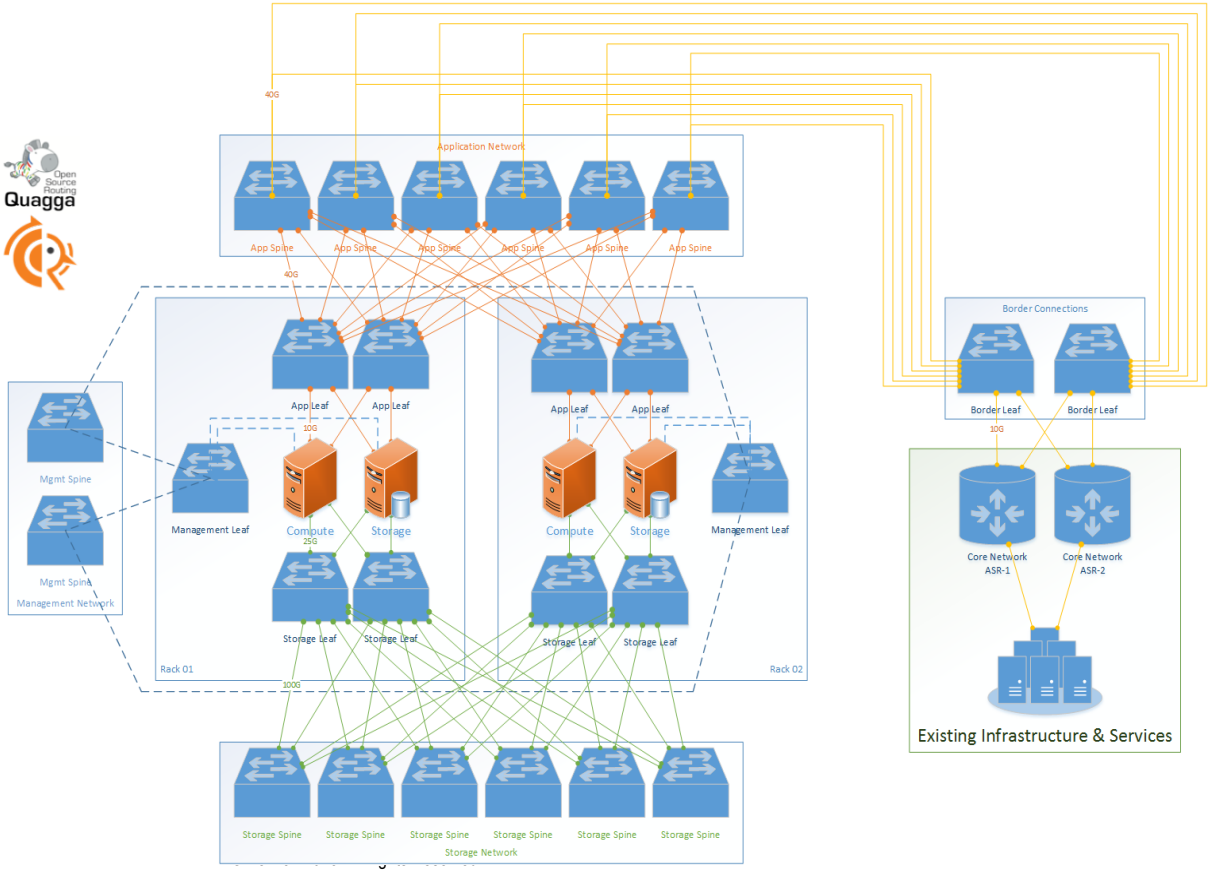

Hardware Platforms

Dell EMC DSS 9000 Servers & Storage

- Open Compute based infrastructure avoids lock-in
- Rack-Level Hyperscale Infrastructure
- Rack-Level Management – IPMI/Redfish
- Flexible Compute & Storage Configurations
- OEM support & reliability, ODM pricing

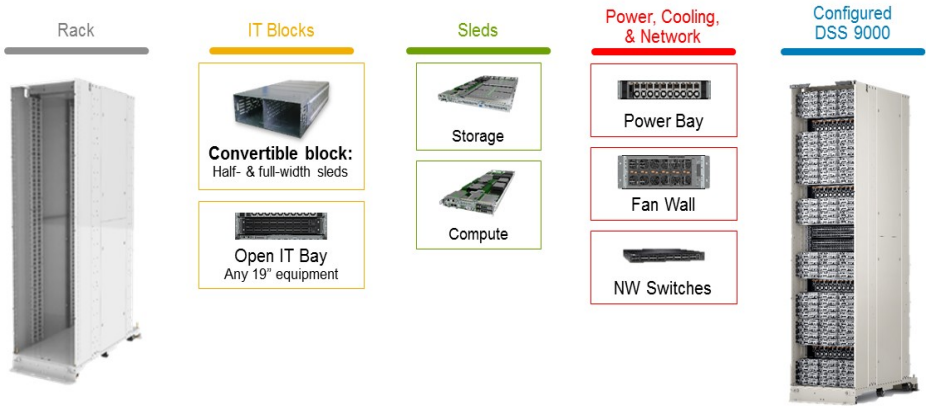

Dell EMC Open Networking Switches

- Swappable Network OS via ONIE
- Latest high performance platform (e.g. Tomahawk 4x25G/100G)
- Commodity silicon avoids vendor lock-in (Z9100 ~= FB Wedge100)
- OEM support & reliability, ODM pricing

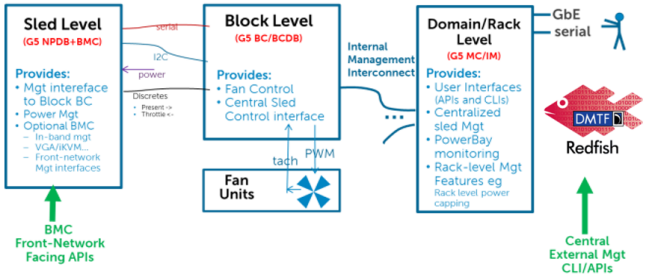

Source: Dell EMC ESI

verizon connect

Infrastructure Stack

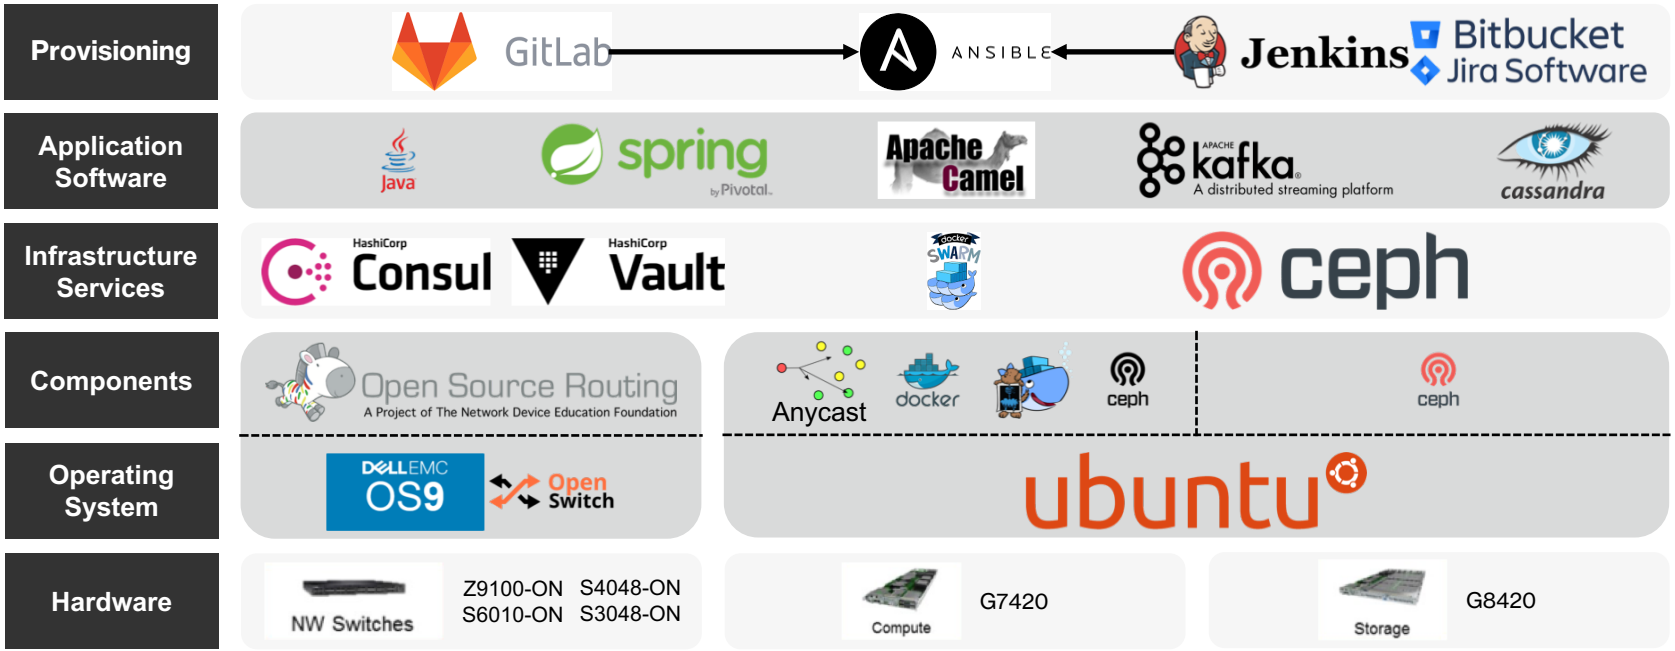

verizon connect

Logical Architecture

Service-by-Service Failover

Service-by-Service Isolation

- VXLAN/MACVLAN Isolation
- Line Rate ACL's in Silicon

Anycast Load Balancing

Ceph Block/Object/File Store

- Persistent Container Storage via REX-Ray

Independent Clusters

- Failure Domain Isolation
- Site-to-Site Replication in App/DB Control

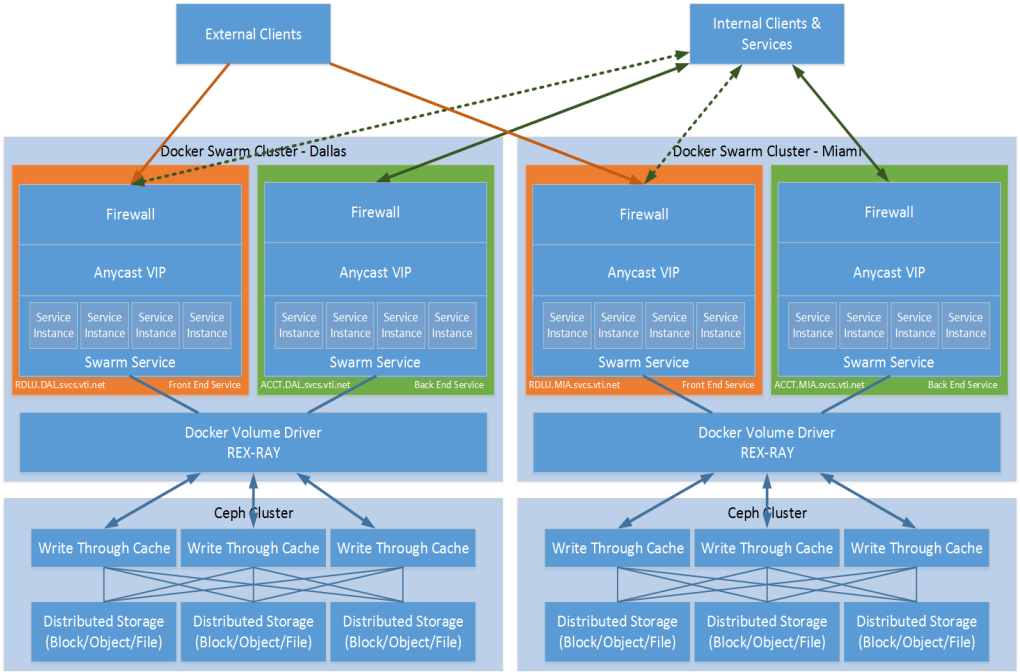

Docker Swarm – Our Perspective (Early 2016)

Decision Factors

- Developer Simplicity & Talent Pool
- Blue/Green Deployment & Rolling Upgrades
- Network, Volume, Logging Drivers
- API Compatibility & Ecosystem Tools
- Windows/Mac/Linux Hosts
- Windows & Linux Containers

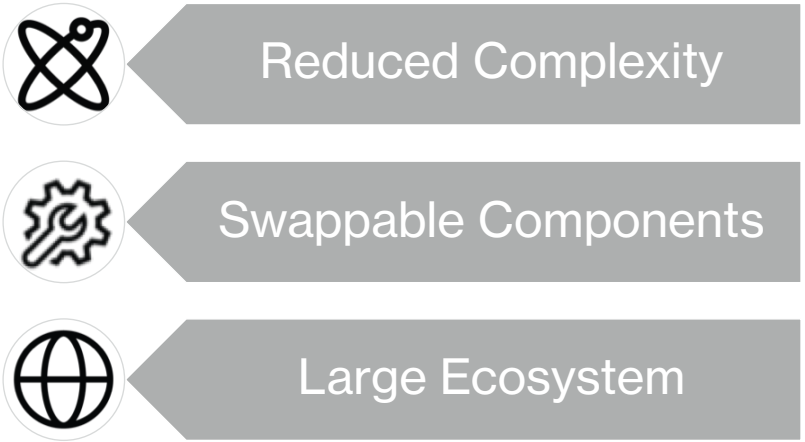

verizon connect

Sounds good…what if it doesn't work?

Swap Components

- **Need cloud?**
  - Swap provisioning to cloud
- **Need Kubernetes?**
  - Enable in Docker or swap it
- **Need etcd?**
  - Add it or swap it

Handover

- **Certify hardware with other VZ platforms**
- **Handover few racks for platform deployment**
- **Convert apps**
- **Validate, update CICD**
- **Deploy in PROD**
- **Convert remaining racks**

Migrate & Repurpose

- **Validate/convert apps and infrastructure services**
  - Service Discovery, Logging, Monitoring, etc, etc
- **Validate, update CICD**
- **Deploy in PROD**
- **Repurpose existing racks**

verizon connect

Automation

How do we keep doing it?

verizon connect

Automation Goals

Simplified

Fully Automated

Easily Repeatable & Predictable

Zero Touch & Zero Downtime

Standardized

Ansible First

Bash/Python Second

Same CI Pipeline, Every Site

End-to-End

Server AND Network

- Configuration AND Provisioning

Public & Private Cloud Adaptable

Automated Maintenance & Healing

verizon connect

CI Pipeline & Deployment Management

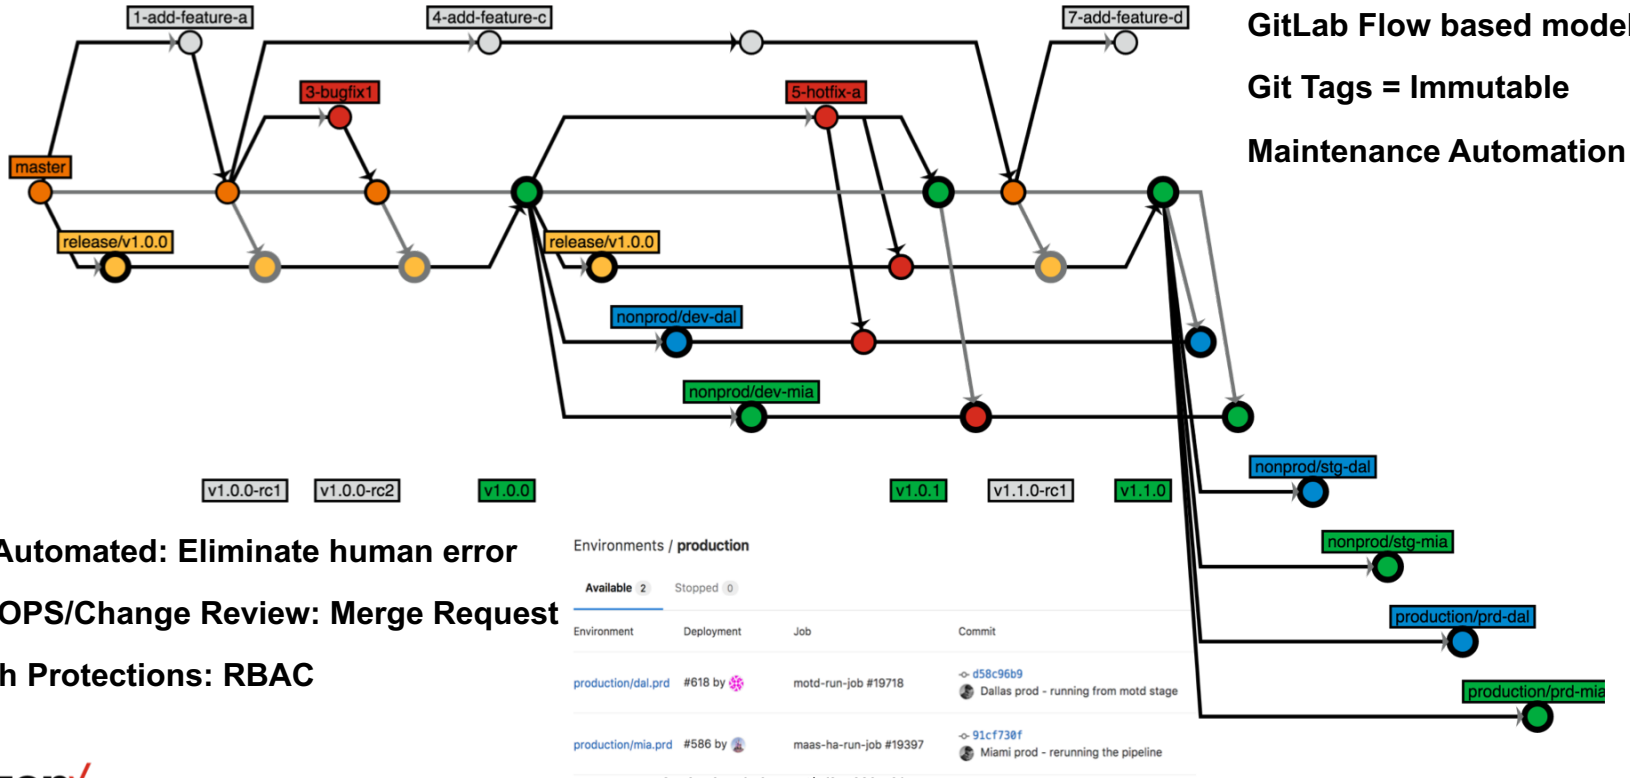

GitLab Flow based model

Git Tags = Immutable

Maintenance Automation

Fully Automated: Eliminate human error

Code/OPS/Change Review: Merge Request

Branch Protections: RBAC

verizon connect

Simple Scalability – Adding Racks

Add rack group (P4)

Update device details

Execute CI Pipeline

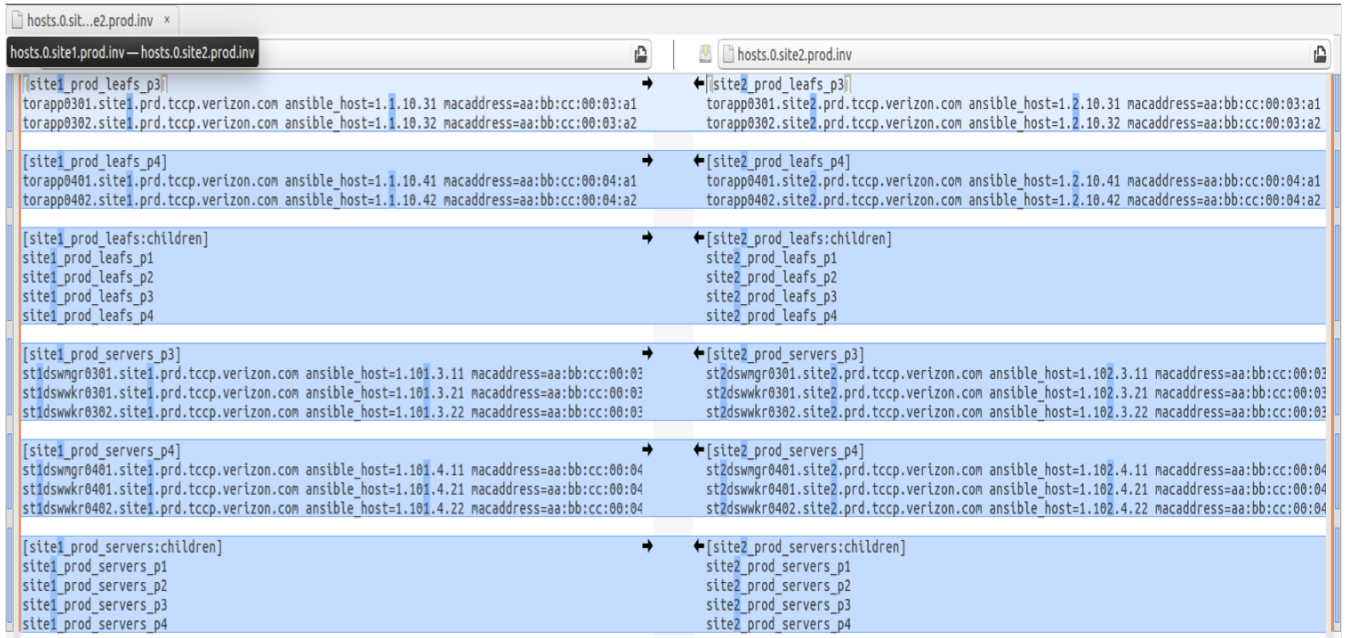

verizon connect

Simple Scalability – Adding Sites

Add datacenter config

Add racks

Execute CI Pipeline

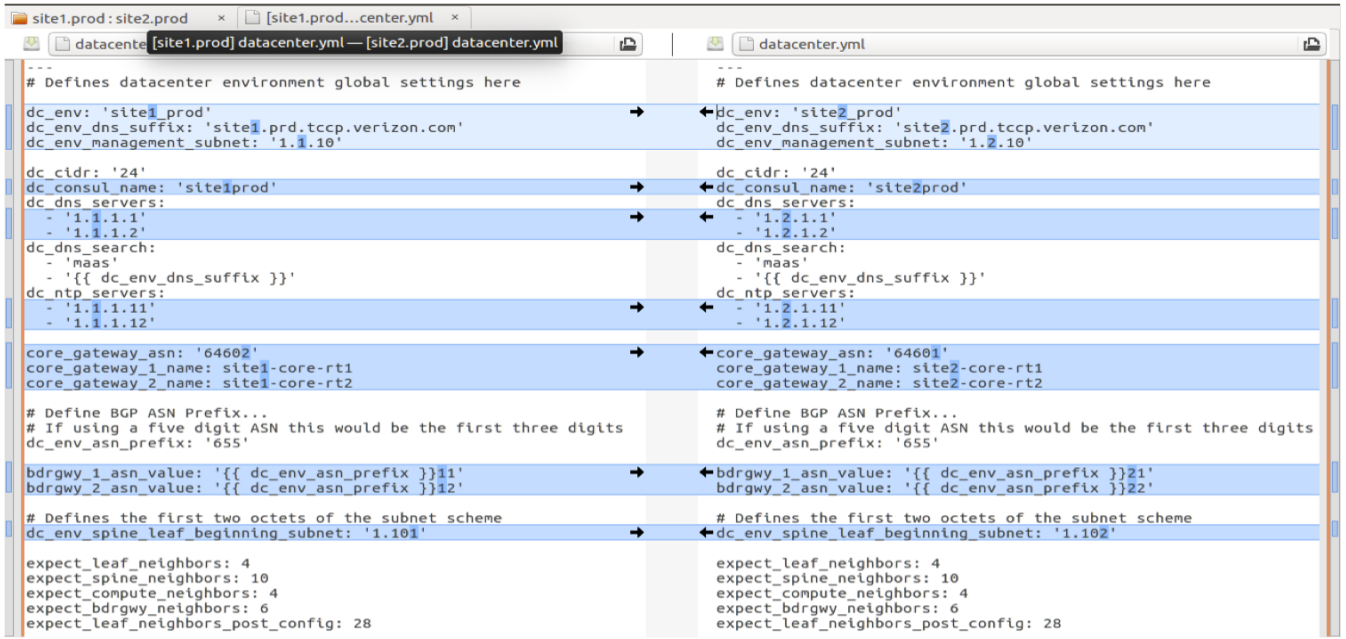

verizon connect

Collaboration & Lessons Learned

What should we have done?

verizon connect

Collaboration

Collaboration

- Issues/Merge Requests, Milestones, and Boards
- GitLab Flow Based Model
- Source Controlled CI
- Reusable repositories

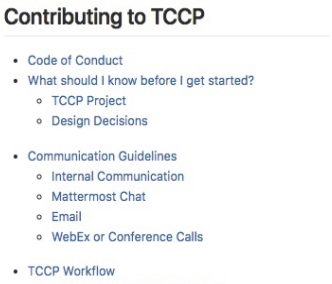

Communication

- Wiki & GitLab Pages
- ChatOps
- CI Analytics

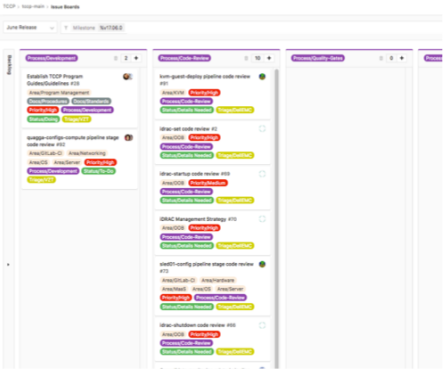

Flexibility

- Single Solution, Multiple Integrations
- Push & Pull Git Mirroring
- Granular Permissions by Team/Group/Project
- Multi-project CI pipelines

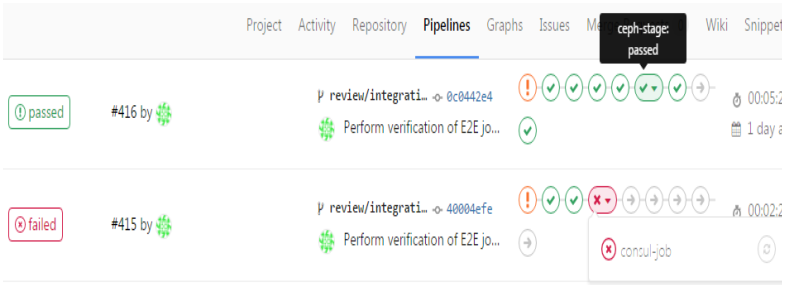

verizon connect

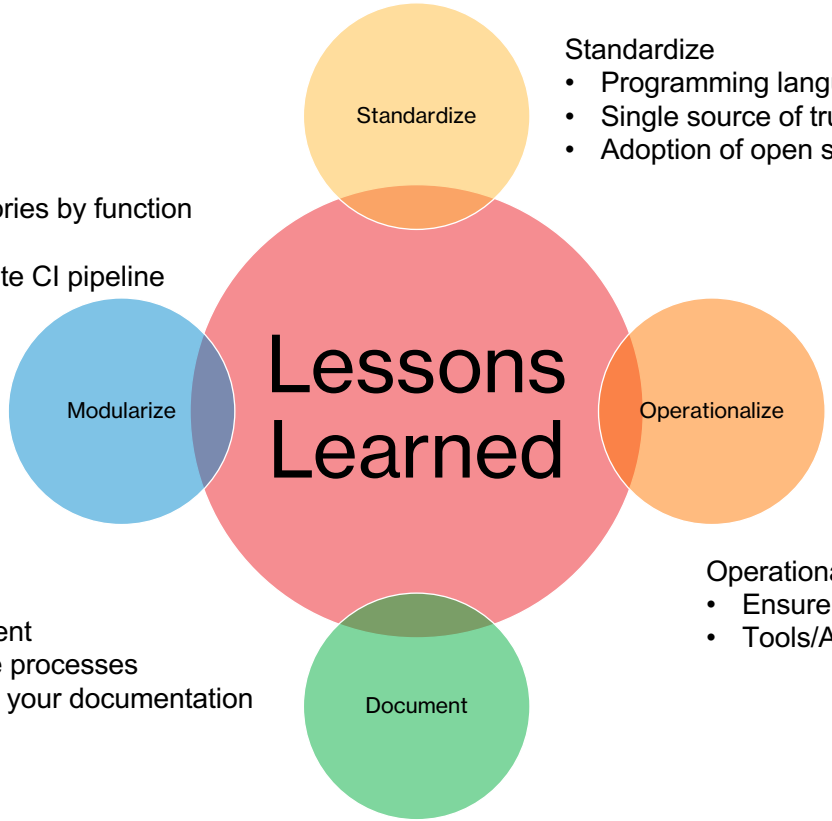

Standardize
- Programming language
- Single source of truth – Ansible inventory
- Adoption of open source roles

Modularize
- Separate repositories by function
- Consolidate roles
- Logically segregate CI pipeline

Operationalize
- Ensure idempotency
- Tools/Alerts/Dashboards in roles

- Document
- Change processes
- Code is your documentation

verizon connect

Q&A

verizon connect

Thank you.

verizon connect

BACKUP

Z9100 Port Mapping & Buffer Zone Utilization

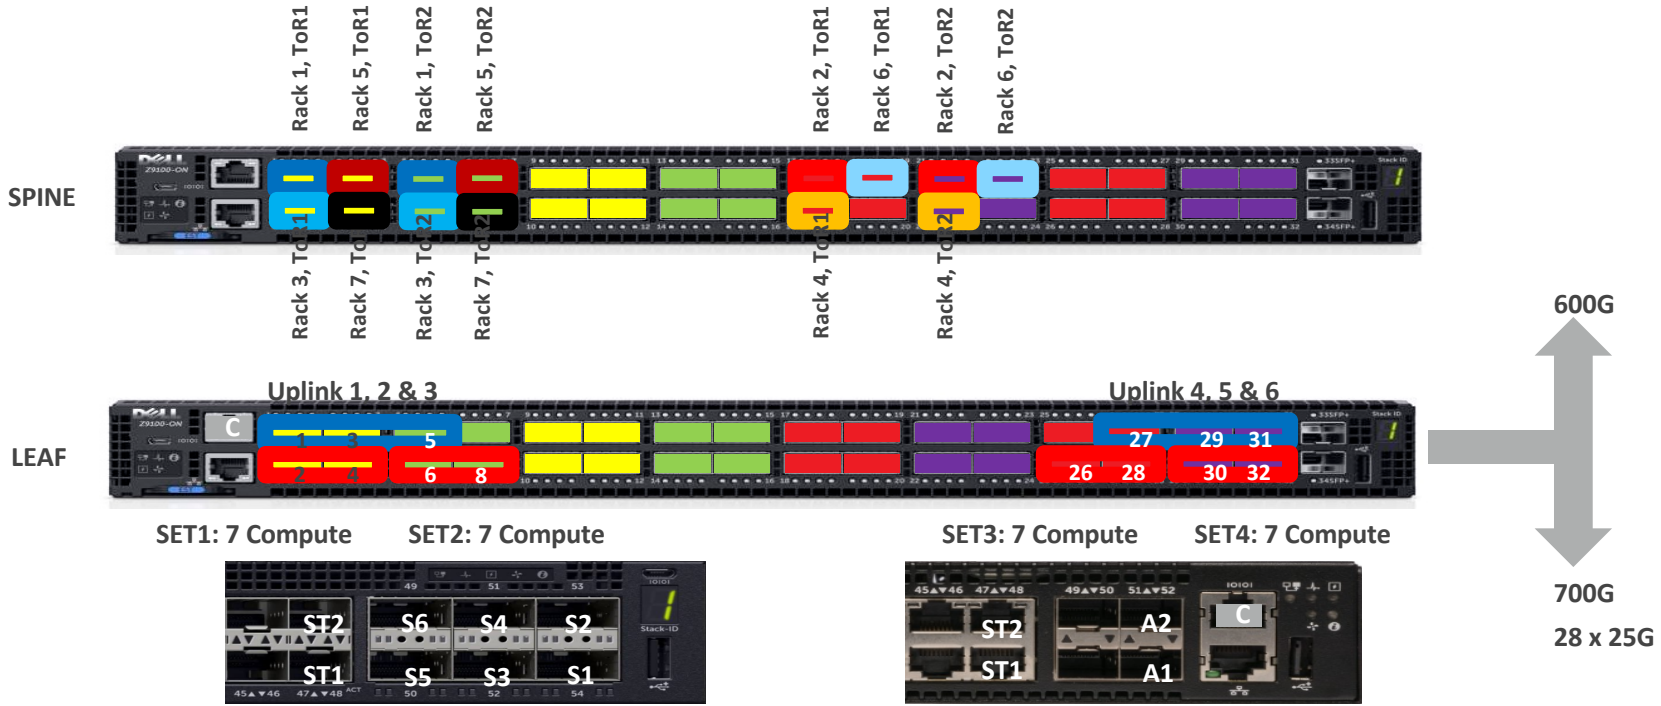

MACVLAN & Anycast is Easy :)

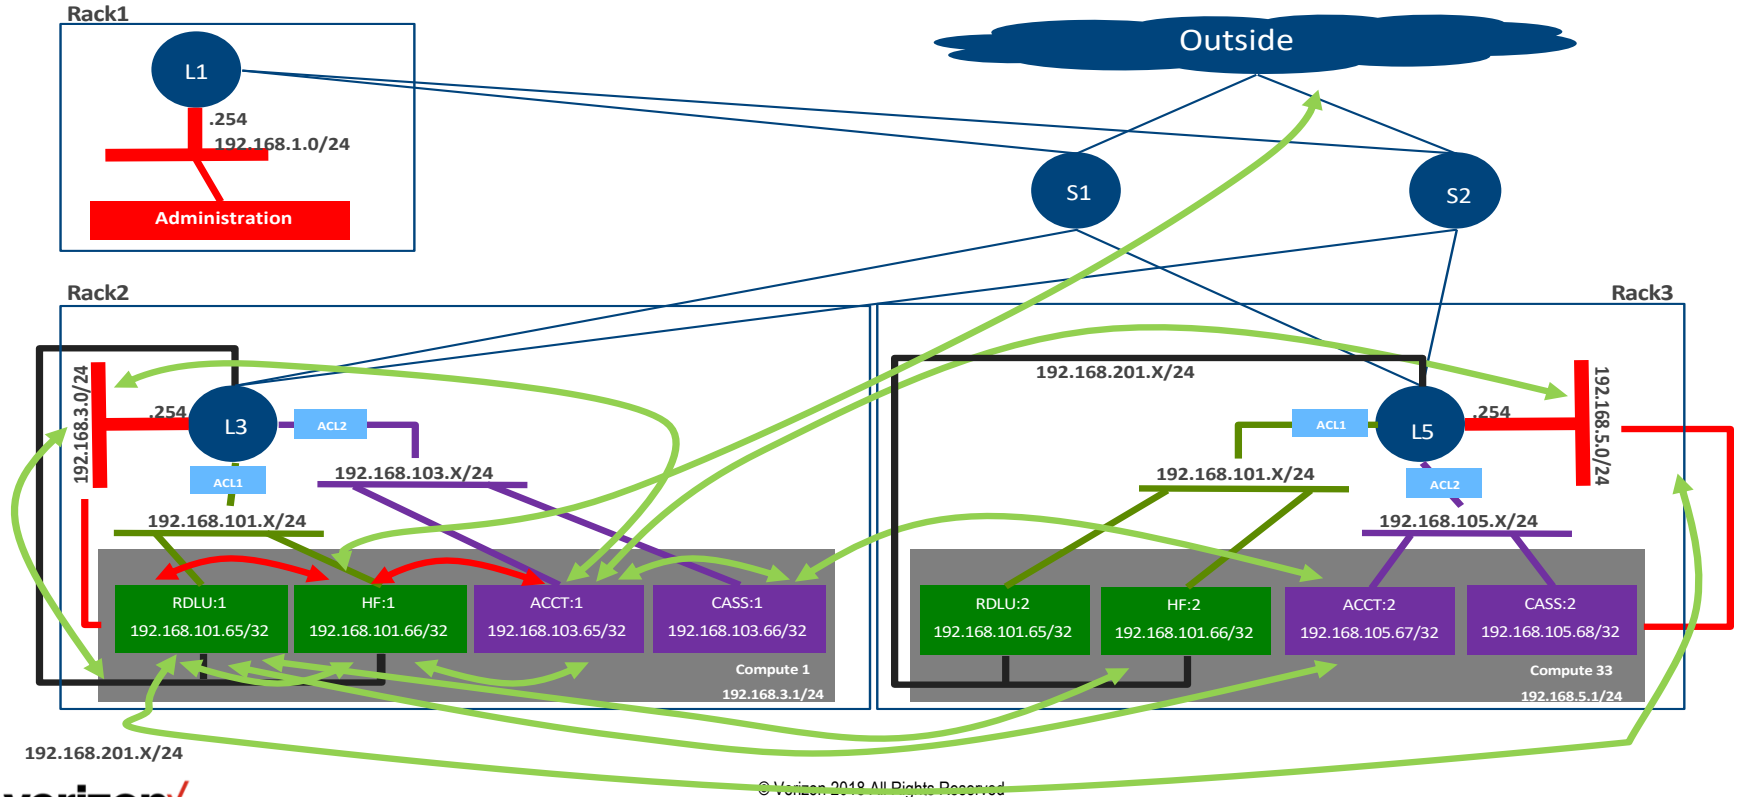

verizon connect