New Data Center Network
DCSS Fall 2019

Eric Brown, Senior Network Architect, NI&S

VIRGINIA TECH™

seen on Facebook, Timothy Holmes Whitmer

Highlights

- 10G ethernet
- Vlans no longer constrained by physical connection
- Network resiliency - available to users
- New cabling infrastructure

Integrated Research & Enterprise Network

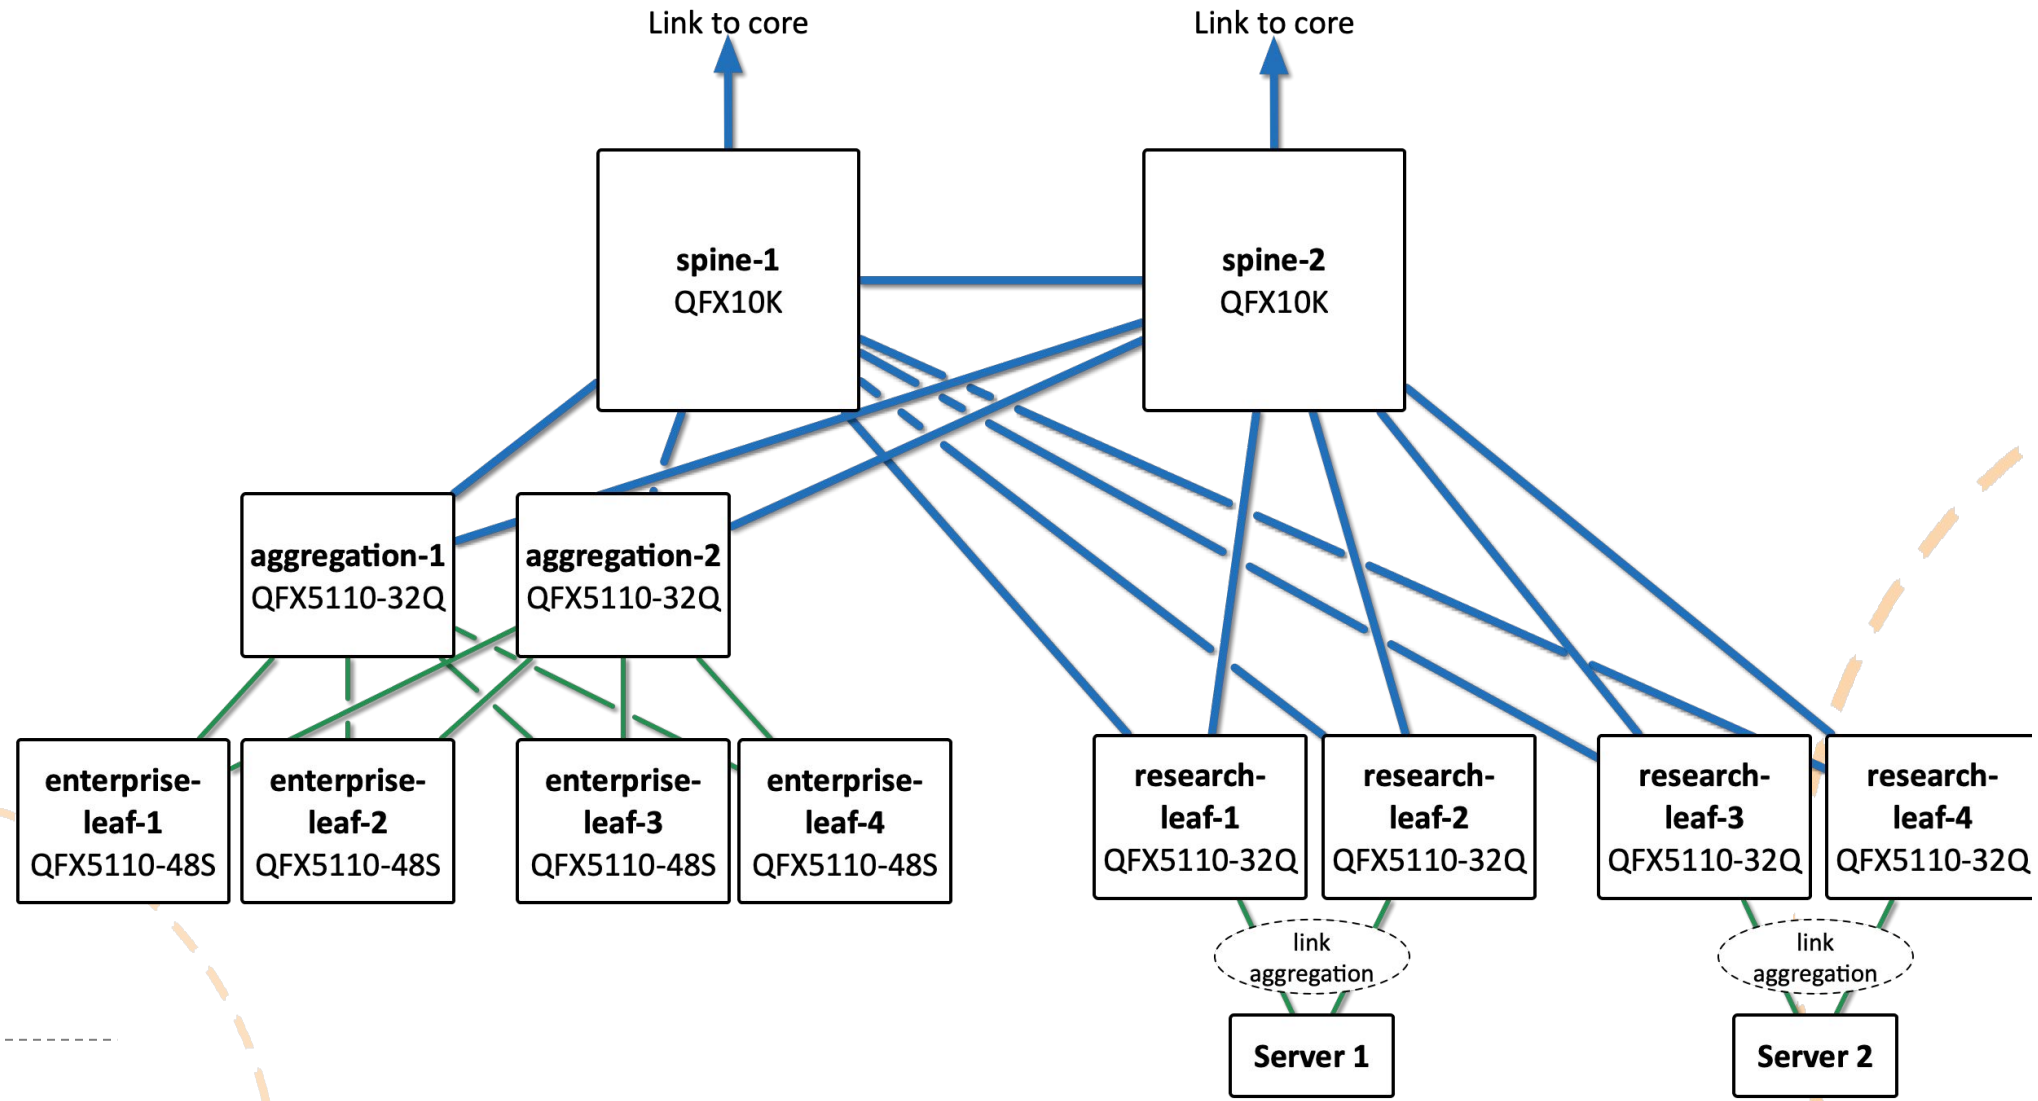

VIRGINIA TECH

Equipment

Spine/Leaf Topology - Horizontal scalability

- Spine: Juniper QFX10K with 30 x 100G
- Enterprise Leaves: Juniper QFX5110 48 x 10G
- Research Leaves: Juniper QFX5110 32 x 40G

Network Architecture

- IP Fabric - all paths equal cost
- N x 100G fabric - research
- N x 40G fabric - enterprise
- EVPN - Combined Ethernet and IP routing for full user connectivity anywhere in the data center
  - VLANS no longer constrained by physical connection
  - EVPN: Ethernet Virtual Private Network

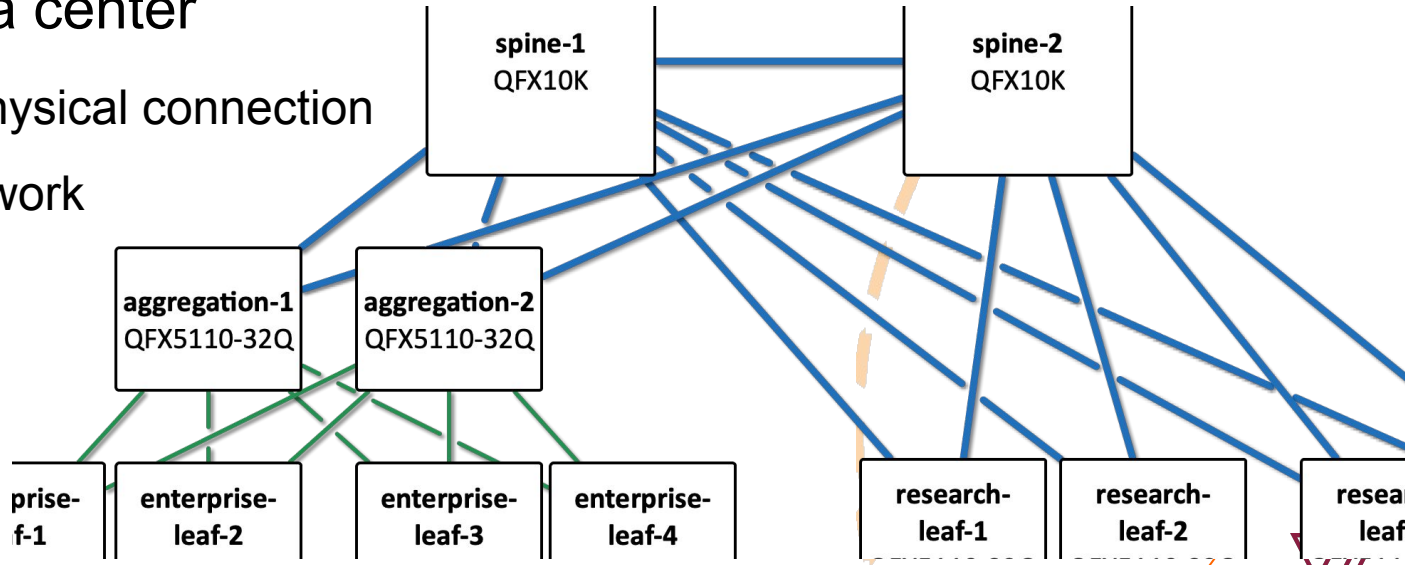

Ethernet Connection Speeds

- General 1/10G support - 100M frowned upon, 10M not available
- 40G research support
- 100G testing

Resilience

- Maintenance without service impact
- Users can take full advantage of this with redundant links

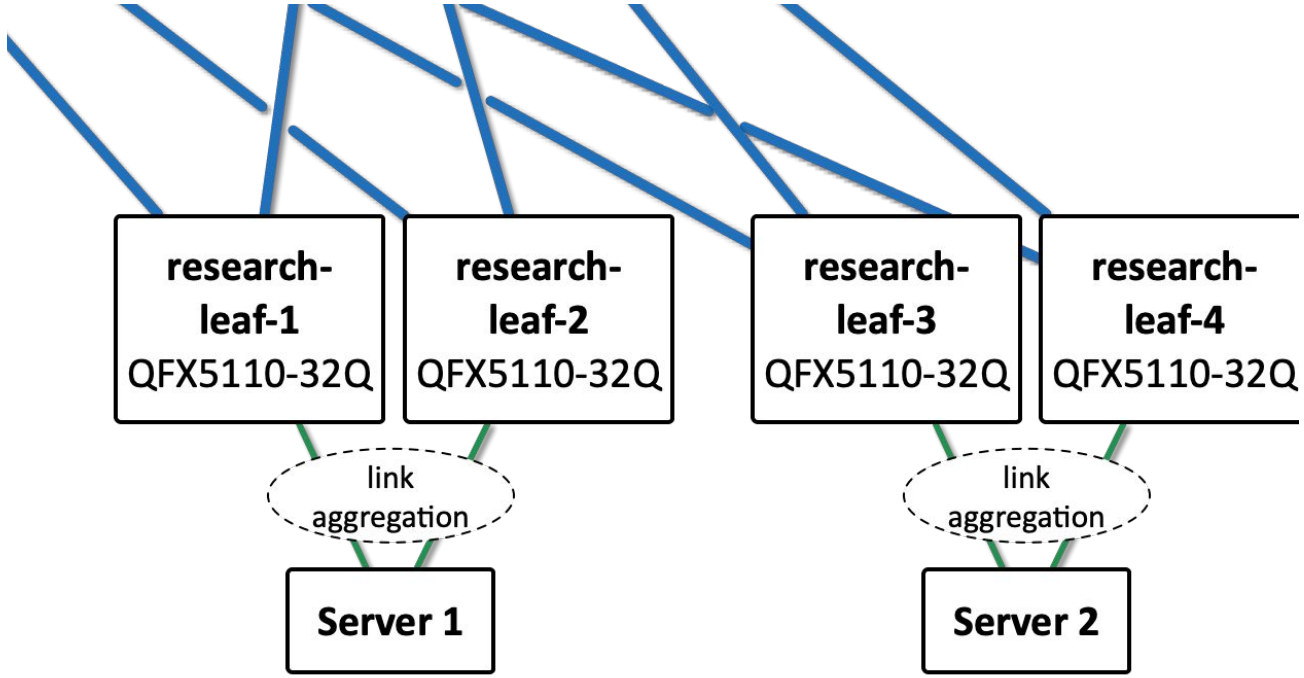

Connectivity

- Only IP connectivity in and out of DC
  - Cannot extend VLANS out of the DC
- VPN support with full policy controlling routing
  - Ex. groups A, B and C can talk to storage but not each other
  - Integrated with campus core VPNs
  - Still an engineered solution

Improved Cable Management

- Overhead cabling
- Zoned switches
- Patch panels

  adjacent to each rack

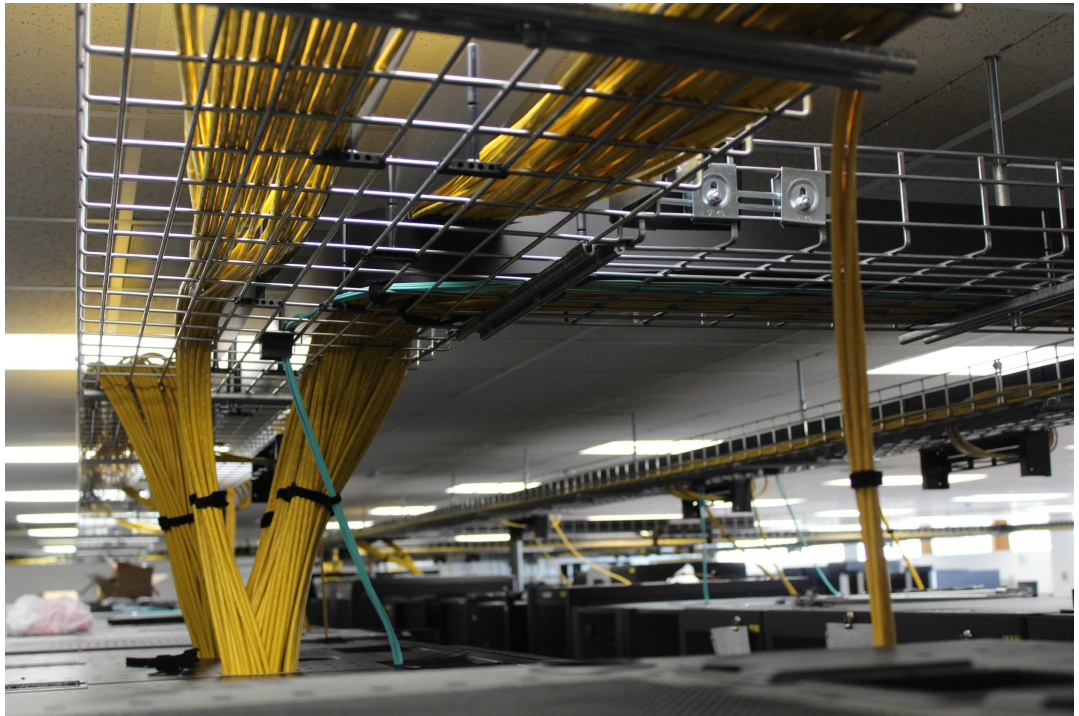

VIRGINIA TECH

Project Phases

- Planning
  - Requirements gathering meeting
  - Technology selection
  - Iterate topology design to meet research needs
  - Testing and validation
  - Purchase

Project Phases

- Deployment and Service Migration
  - Agile deployment project
    - deploy network, develop config automation, manage migration
    - communication overhead, but returns in quality, velocity and problem management
- Issues and Learning Opportunities
  - Virtual MAC address for router
  - Routing instability due to network scans
  - Smooth migration of connections across leaves inadvisable
  - Certain mis-configs can be harmful, mitigate with automation

Future: Automation and Orchestration

- We need your help!
- Looking for opportunities to collaborate with users on solutions

Future Development - A network is more than moving packets

- User configured security and segmentation
- Support bare-metal provisioning
- Load balancing service - automation and re-architecture
- DNS automation/orchestration
- Network visibility
- IPv6 first data center

Thank you!

Shout out to the team

Steve Lee

Gary Hess

Ajinkya Fotedar

Pranav Baitule

Nandan Sadineni

Shane Kemp

Brian Jones

Sara Morris

VIRGINIA TECH

Questions?

Architecture Details

- EVPN
  - MAC address learning happens in the control plan as opposed to the traditional flood/learn mechanism in the data plane
  - Ability to stretch broadcast domains across multiple racks
  - Support for active-active multi-horned LAG connections using the Ethernet Segment Identifier (ESI) object
  - Extension to the well known protocol used in the network core and internet
  - Supports multiple tunneling technologies (VxLAN in DCN)
  - Support for hitless upgrades for any emergency maintenance

Architecture Details

- Full IP Fabrics
  - No Spanning Tree Protocol
  - Eliminates escalation of broadcast domain failures into network meltdown
  - Equal-cost multipath and load balancing
  - 100G links and large buffers support mice (enterprise) and elephant (research) flows

Architecture Details

- Separation of (overlay) user traffic from the (underlay) traffic that is used for maintaining fabric's reliability
- A spine-leaf design helps ease future growth with horizontal scaling
- Automated fabric deployment, migrations and service provisioning