


Virtual Topology Service


GEC 18



BARNSTORMER
S O F T W O R K S


Agenda

- Audience
- Goals
- Architecture / How it Works
- How You Can Help!




Goals

- More Flexible Topologies
- DataPath Programmability
- Increased Visibility
- Standard Tools / Management Support
- Enable High Quality Research




Flexible Topologies

- GENI WAN Is Your Oyster
- Allow at least a full mesh of rack edges
 - Behind-the-scenes effort to avoid massive path oversubscription
- Some decision points in the “core”
 - More limited options, but useful in some topologies




Flexible Topologies

- Use any of a wide variety of DataPath implementations
 - OpenFlow (OVS, IVS, Linc, xDPD, etc.)
 - Cisco OnePK
 - Click
- Integrate Middleboxes as first-class members of your topology
 - BigIP, Checkpoint, etc.




Programmability

- Can optionally use an entirely custom datapath implementation
- Create and deploy extensions to existing datapath implementations, where possible (OF, etc.)
- Arbitrary dataplane pipelines
- Modern SDN control plane support
 - Distributed / HA / Failover




Increased Visibility

- All links in your topology can be tapped / monitored in isolation from other slices
- Full access to your datapath for debugging at code level
- No control plane proxy (FlowVisor, etc.)
- Access to underlying infrastructure utilization information across the time period of your slice




Standard Tools Support

- Direct access to monitoring data from your datapaths
 - sFlow/NetFlow
 - SNMP
- Can use NCM/SCM / NMS tools
 - Zenoss / Nagios / Hyperic
 - NetDisco / HyperGlance




Enabling Research

- IPv6, MPLS, FC
- End-to-End QoS, Flow Control
- Latest versions of OpenFlow
- Non-OpenFlow SDN Platforms
- Non-SDN Platforms




How It Works

- L2 Pseudowires and Virtual Machines
 - You can set this up now, but there's a lot of heavy lifting
- Some new core hardware coming online that has better slice isolation support





How It Works

- Aggregate managers and tools set up bridges and paths for you
- Paths created via:
 - Existing OpenFlow Core
 - DCN Aggregates (OSCARS, AL2S)
- Reasonable efforts will be made to provide backup paths and reduce primary path contention





Network Architecture




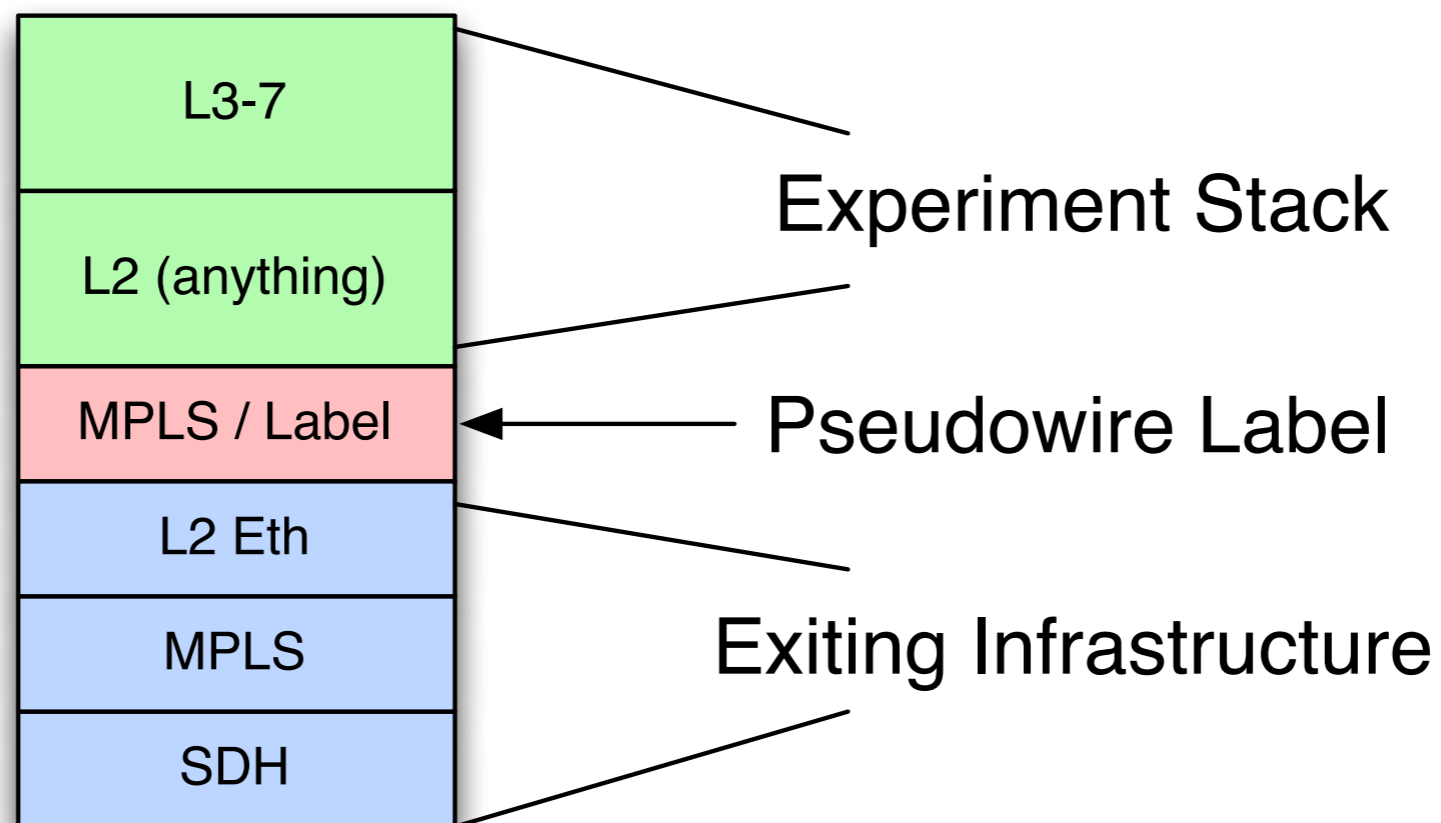
Network Architecture

- Experimenter defines the network graph that they want created using any vertex (rack, core resource) available in the system




Pseudowires

- Typically MPLS edge-to-edge
- Opaque to the core switching devices




Pseudowires

- We unroll the pseudowires at each experiment interface so your datapath implementation and hosts are directly exposed to your L2 of choice
- If your experiment uses bare metal there are some (mostly obvious) limitations




GENI Pieces

- Aggregate Manager at each site (rack) that supports pseudowire endpoints
- Request specs:
 - Specify **simplex** pseudowires
 - Datapath images to connect them to
- Labels are only locally unique
 - Avoids nasty coordination problems



Limitations

- You must (should) have a datapath at each site where you have compute resources
- You cannot (currently) specify the L1 path you want the pseudowire service to use
- You can't actually reserve any resources right now



Tool Writers!

- Topology creation / resource coordination will be a pain to do by hand (much like now)
- Documentation and rspec schemas will be available after GEC
- Beta deployment hopefully available before GEC19

