
SDN SANDBOXES

An approach for GENI Mesoscale



Current Issues - Data Plane

- * Physical hardware has limitations
 - * Not all matches can be done at line rate
 - * Very few actions are available to experimenters
 - * Extremely limited table size
- * Firmware upgrades are long-tail
 - * Even switches whose hardware could support full OpenFlow 1.0 won't have firmware support for another 6-18 months

Current Issues - Control Plane

- * Lack of rate-limiting to prevent switch CPU starvation
- * OpenFlow 1.0 is the only supported SDN platform
 - * ONF spec velocity far outpaces hardware upgrades
 - * Experimenters don't have access to competing technologies
 - * ForCES
- * FlowVisor is a single point of failure and fast-path contention

Current Issues - Architecture

- * The infrastructure is not decoupled from the experiment
- * Experiments can take down infrastructure
- * Experiments and the infrastructure contend for resources
- * Constrained label space
 - * VLAN tags are not available to experimenters
 - * Rewrite / Translation not supported at the edge

Optimal Solution

- * Isolated L1
 - * via nationwide long-haul DWDM
- * Provided by?



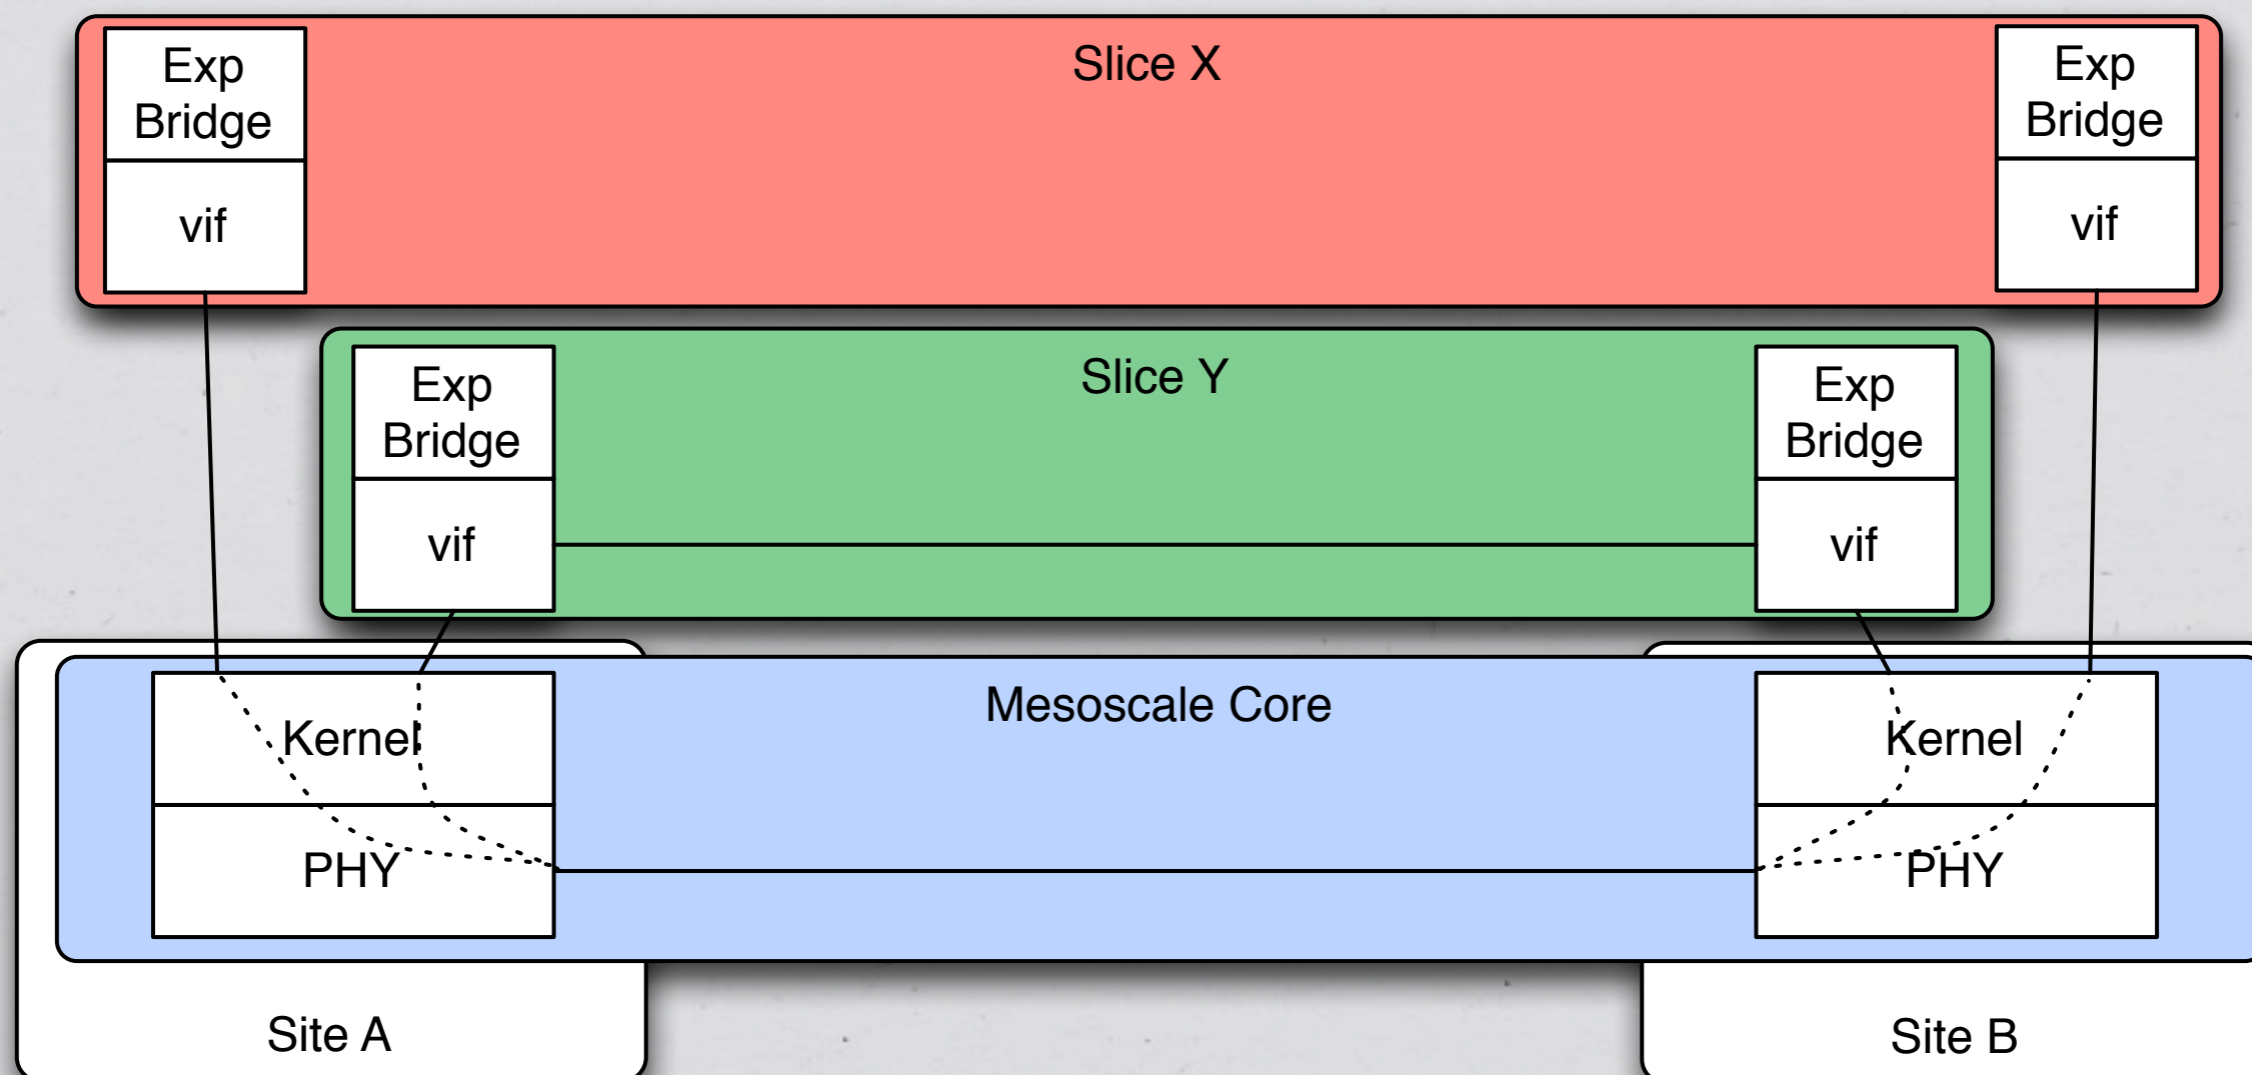
Practical Solution

- * Build a substrate on top of the Core
 - * Core Mesoscale becomes a transit provider
 - * (ala I2, NLR, etc.)
 - * Provide isolated L2 on top of existing L2(.5)
 - * Already the case today with MPLS, etc.
- * Run the Core network as true infrastructure

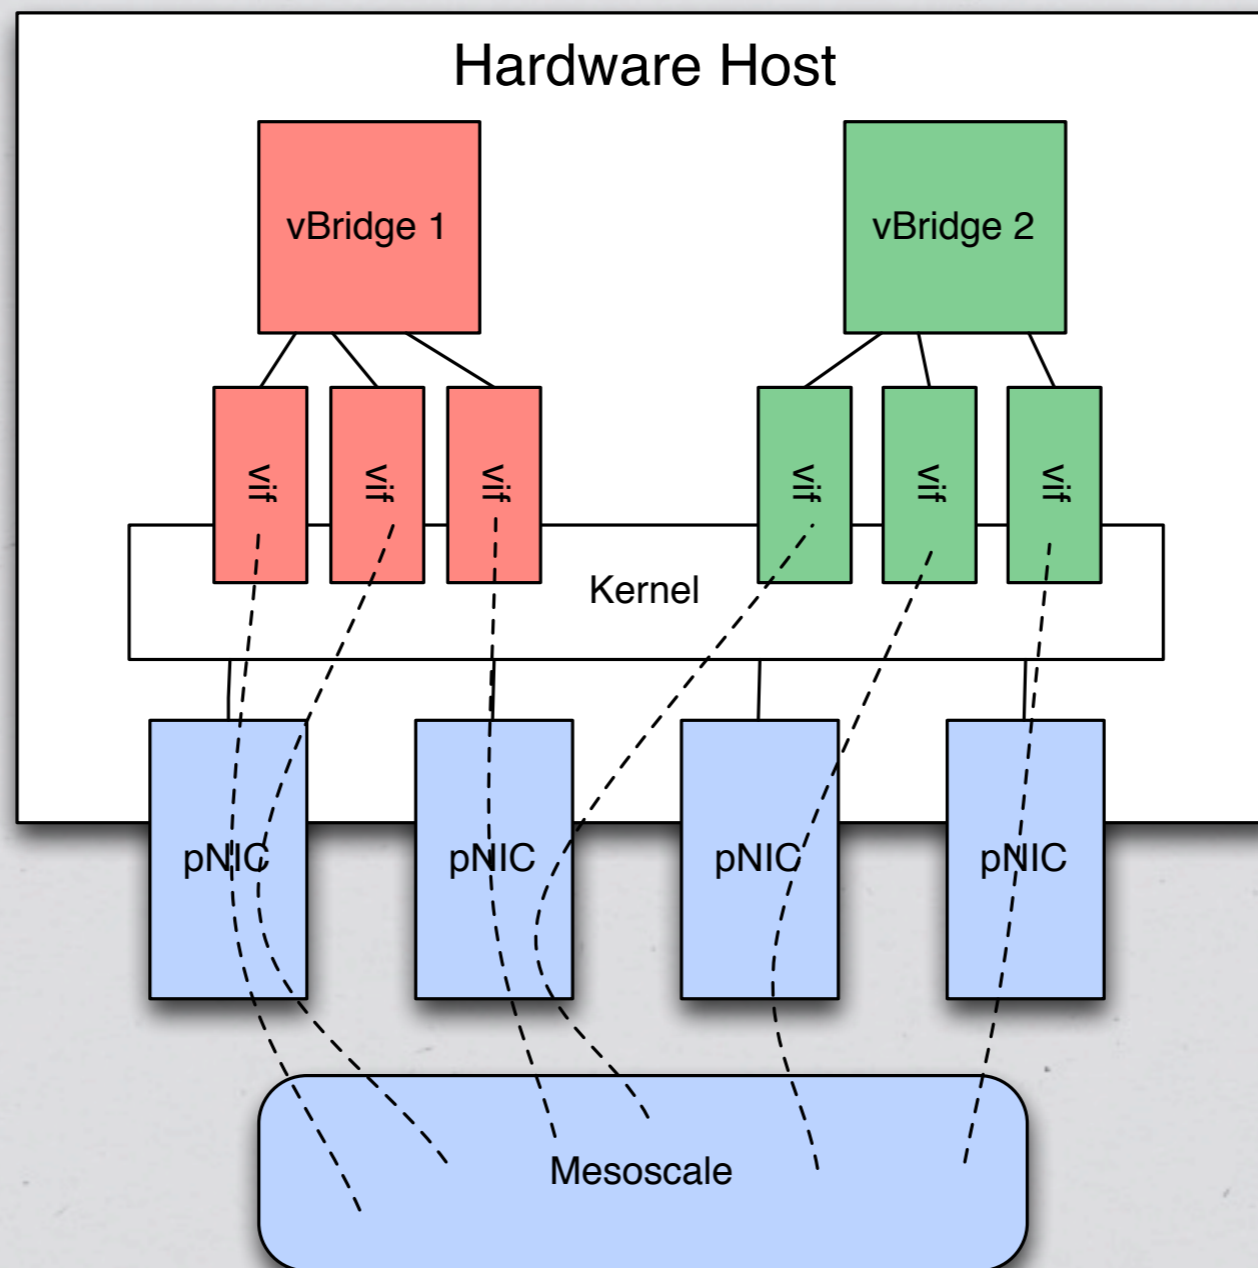
How?

- * Provide isolated virtual networks to experiments
 - * Virtual Links and Dynamic Datapaths
 - * Mesoscale provides transit and termination
 - * Core / Regional / Campuses provide as many datapath instantiation points as possible
- * PCs at as many POPs as possible
 - * A variety of datapath “images” are provided

Experiment Isolation



Experiment Isolation



Advantages

- * Reserve topology, not flowspace
- * Significantly less (100x+) flow table contention
- * Complete L2/L3/L4 space available to each experiment
- * Your datapath doesn't have to be OpenFlow
- * Provides for IPv6 experimentation (and many others)
- * More opportunities / points of instrumentation