

# DTunnels Year 1 Summary

Nick Feamster

# Overview

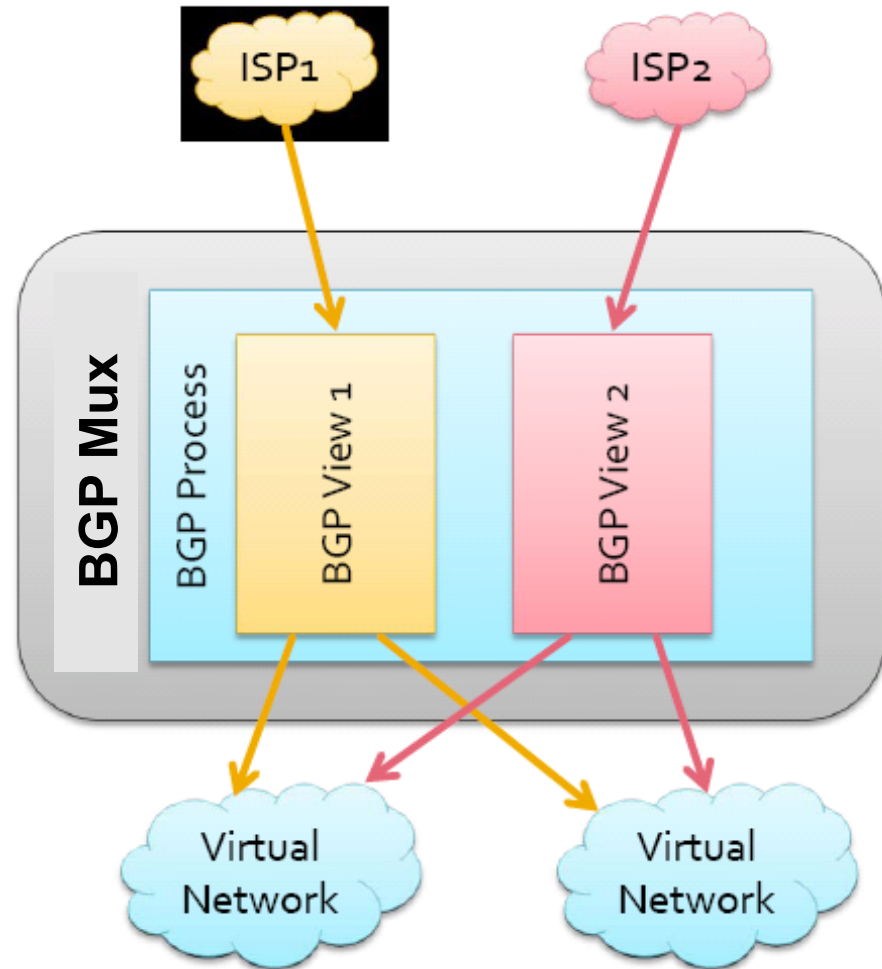
- Two pieces
  - DTunnels: Mechanism for creating appearance of layer 2 links between virtual nodes
  - BGP Mux: Mechanism for connecting downstream virtual network to upstream connectivity
- Two students: Yogesh Mundada (DTunnels) and Valas Valancius (BGP Mux)

# Design Requirements

- **Session transparency:** BGP updates should appear as they would with direct connection
- **Session stability:** Upstreams should not see transient behavior
- **Isolation:** Individual networks should be able to set their own policies, forward independently, etc.
- **Scalability:** Mux should support many networks

# BGP Mux

- Separate BGP view for each upstream ISP
- Private AS numbers are removed before propagation
- “local-as” function for client connections
- Small advertisement interval
- Unchanged BGP attributes



# Progress in Year 1

- DTunnels
  - Kernel patches to create “Ethernet GRE” tunnels
  - Interface for specifying topology in XML and instantiating topology in OpenVZ (ProtoGENI nodes will run)
- BGP Mux
  - Design and implementation of control plane
  - Deployment in three locations: GT, Wisc, PSG

# Challenges in Year 1

- DTunnels
  - Implementation of kernel patch for EGRE tunnels (done)
- BGP Mux
  - Mostly logistical problems
  - Acquisition of AS number, IP address space, mapping of the IP space to the AS number in shared whois and the IRR
  - Data plane integration

# Integration Progress (Spiral 1)

- DTunnels
  - Instantiation of topology from an XML spec that resembles an RSpec
  - Some initial discussions about the commonalities between the ProtoGENI RSpec and our spec
- BGP Mux
  - Installation of Mux nodes on same subnets as ProtoGENI nodes
  - No RSpec integration yet: need aggregate manager

# Internet2 Connection Progress

- BGP Mux is deployed at three sites: Georgia Tech, Wisconsin, PSNet
  - First two sites: directly connected to Internet2
  - Route advertisements are being accepted on Internet2 connections
  - Advertisements filtered on commodity links due to the rwhois/IRR problems
- DTunnels
  - Not really applicable, but we do have the option of provisioning between these sites



# Plans for Experiments

- First experiment: “NameCast”  
(Jennifer Rexford)
  - DNS resolvers are replicated at multiple sites and advertised on a common IP prefix
  - Service advertises and withdraws BGP routes to control how traffic reaches the service
  - Deployment in progress at Georgia Tech
- Next experiment: Something that requires virtual networks/tunneling behind the Mux

# Plans for Spiral 2

- Deployment of experiments on BGP Mux and DTunnels
- Integration of aggregate manager with ProtoGENI (front end)
  - Automation of BGP Mux setup from RSpec
- Integration with other control frameworks