

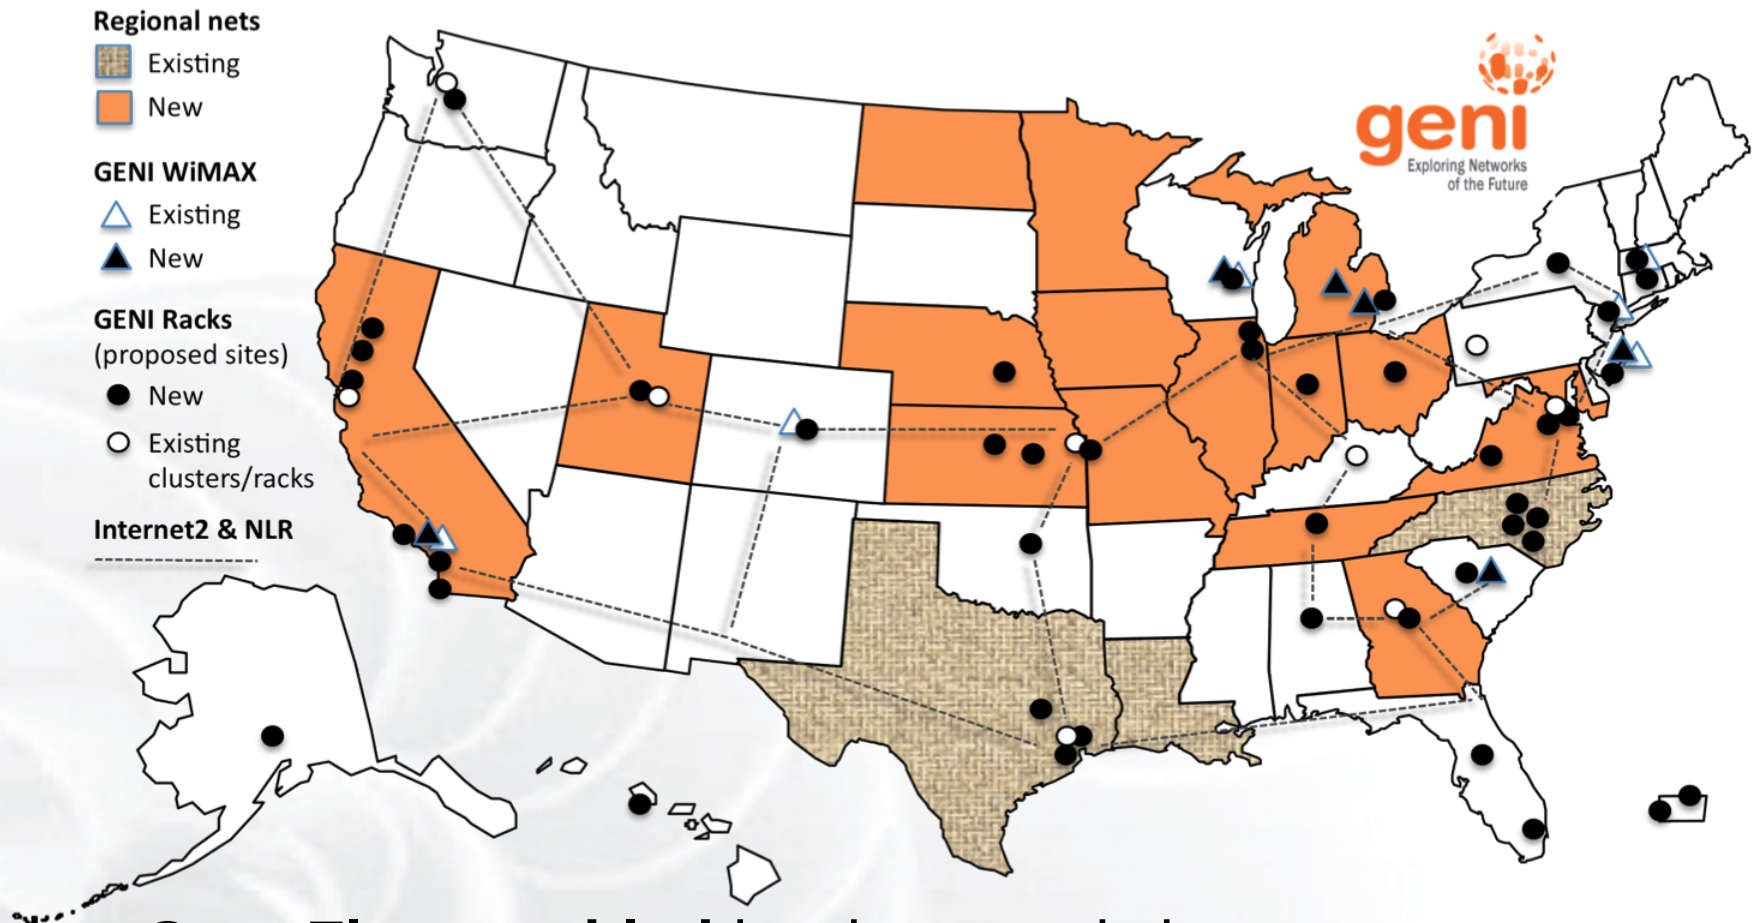
Train The Admin: OpenFlow and FOAM

Josh Smift, GPO
Atlanta GA, 2014-03-18
www.geni.net

- OpenFlow is an API that switches implement
 - Version 1.0 is most common on current GENI switches
 - Version 1.4 is the most recent version
- OF is an example of Software-Defined Networking
- Traditionally, forwarding logic is in the switch
 - The firmware might give you some control (e.g. ACLs)
 - Configuration/programming options very limited
- OpenFlow moves that to an external controller
 - The controller manipulates the switch's flowtable
 - Traffic that matches flow rules is forwarded at line speed
 - Unknown traffic is sent to the controller

- Switches are *programmable*, not just *configurable*
- Innovation is fast and easy
 - Write controllers in any language, run on any system
 - No need for switch-specific firmware
- Controllers can match on various layers
 - Layer 1 (switch port)
 - Layer 2 (VLAN ID, MAC src/dst, ethertype)
 - Layer 3 (IP src/dst/protocol)
 - Layer 4 (TCP (or other protocol) src/dst)

- An OpenFlow switch has one or more *datapaths*
- Each datapath is essentially a virtual switch
- Every datapath has a unique identifier (*DPID*)
- Types of OpenFlow switches:
 - VLAN-based hybrid: Multiple DPIDs, based on VLANs
 - Port-based hybrid: One or more DPIDs, split by port
 - Pure OpenFlow: One DPID, no traditional switch config
 - Virtual switches (e.g. OVS)



OpenFlow-enabled hardware switch at:

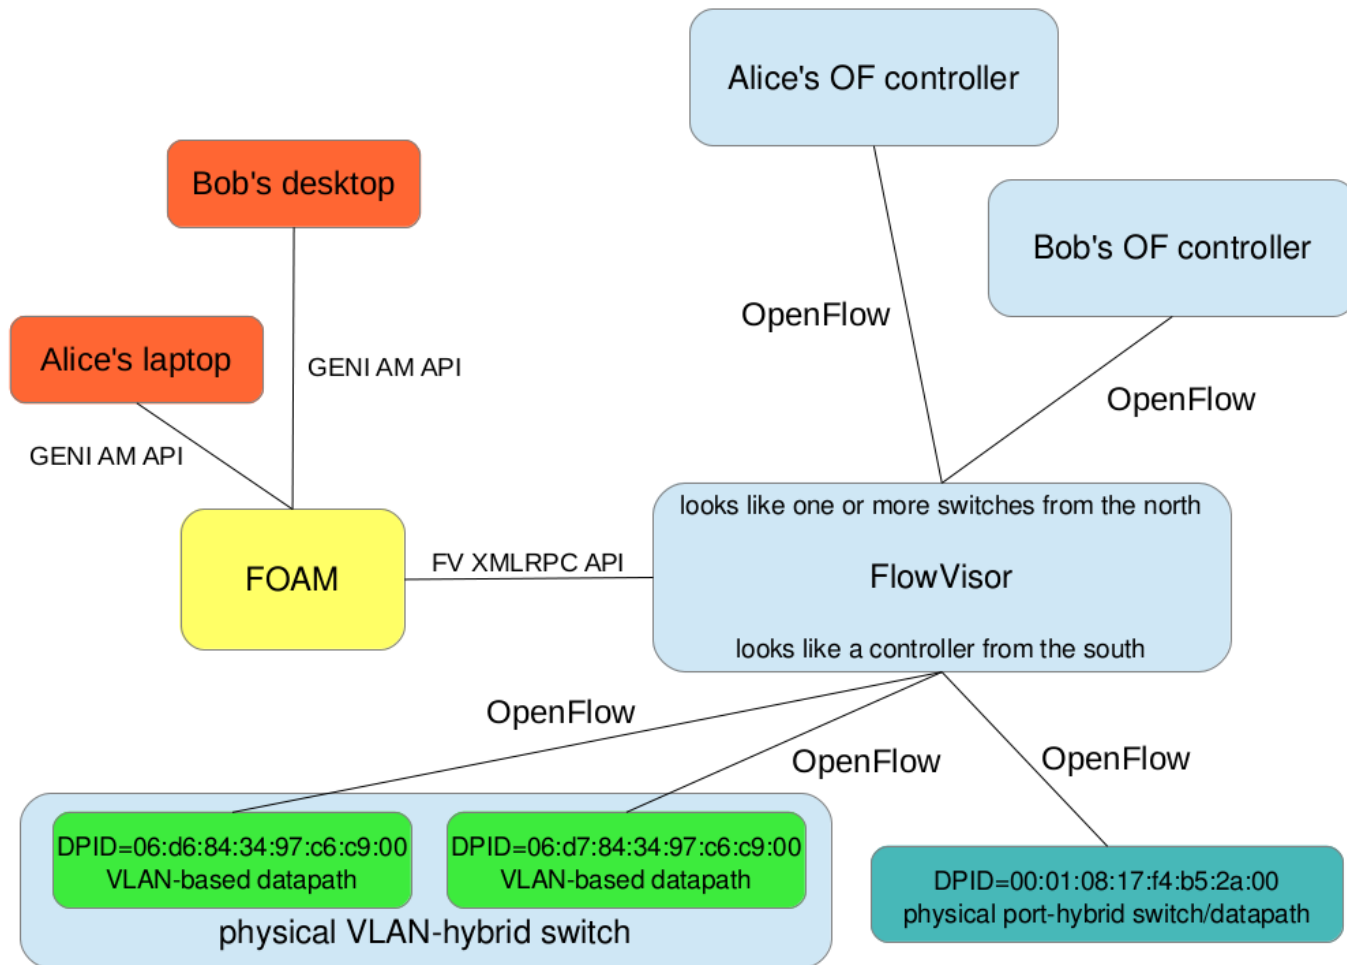
- Each GENI Rack
- Backbone and regional networks

- Indiana University runs an OF class for admins
- Offered periodically at conferences and meetings
 - NANOG, et al
- For more info:
 - Contact Steven Wallace of IU: ssw@iu.edu
 - Slides from the class: <http://goo.gl/sXV6FV>
- GEC also has non-admin-specific OF tutorials

- FOAM is an OpenFlow aggregate manager
 - Experimenters can get resources via the GENI AM API
 - FOAM uses FlowVisor to allocate OpenFlow resources
 - ...much like ProtoGENI uses Xen to allocate VMs
- <http://groups.geni.net/geni/wiki/OpenFlow/FOAM/AdminIntro>
 - All the material from this part of the presentation
 - Lots more details and links to other useful resources
 - Recommended reading for new FOAM admins
- gpo-infra@geni.net
 - The GPO ops/infrastructure group
 - A fine place for any questions/comments/etc

- A physical switch has one or more *datapaths*
- A datapath is essentially a virtual switch
- Each datapath has a unique identifier (*DPID*)
- FlowVisor slices traffic into *flowspace*
 - Each experimenter has their own flowspace
 - FlowVisor provides separation between them
- FlowVisor is completely controlled by FOAM
- FOAM and FV can run on one host, or two
 - InstaGENI racks have two VMs on the control node
 - ExoGENI racks run both on the head node

FOAM architecture diagram



- In other aggregates, all requests are approved
- In FOAM, flowspace requests can overlap
 - For example, Alice can request IP subnet 10.42.0.0/16
 - What if Bob then requests 10.0.0.0/8 or 10.42.15.0/24 ?
- All requests used to require manual admin approval
- Nowadays, requests are auto-approved if no overlap
- If a new request does overlap existing slivers:
 - The request will be marked as “pending”
 - You'll get mail from a nightly cron job about it
 - The experimenter will usually need to delete and try again
 - Ask gpo-infra@geni.net with any doubts/questions/etc

- FOAM sends mail immediately when:
 - Experimenters act on slivers (create/renew/delete/etc)
 - Slivers are approved (automatically or by hand)
 - Slivers are disabled (by the admin)
 - NOT when a sliver is created but not auto-approved
- FOAM cron jobs send mail when:
 - Slivers expire (checks every six hours)
 - There are pending slivers that need approval (nightly)

- Pending slivers require action
 - Experimenter will usually need to delete and try again
 - Or the existing sliver may have reserved too much
 - Savvy experimenters will know something is amiss
 - Naive experimenters may not know what's wrong
- Make sure to keep an eye out for that cron mail

```
From: FOAM at BBN <foam-admin@gpolab.bbn.com>  
To: foam-admin@gpolab.bbn.com  
Subject: [gpo-ops] FOAM Pending Queue  
Date: Thu, 12 Sep 2013 09:52:02 -0400 (EDT)
```

(This is an automated message from FOAM.)

The following slivers at foam.gpolab.bbn.com are pending admin approval:

```
Sliver URN: urn:publicid:IDN+ch.geni.net:JBS+slice+jbsfoamtest:a4ee89ac-3f7b-46fa-a78f-  
bf24a22f9474  
User: jbs@bbn.com [urn:publicid:IDN+ch.geni.net+user+jbs]
```

- Other routine messages don't require action

```
From: FOAM at BBN <foam-admin@gpolab.bbn.com>  
To: foam-admin@gpolab.bbn.com  
Subject: FOAM sliver created  
Date: Fri, 10 Jan 2014 13:29:27 -0500 (EST)
```

(This is an automated message from FOAM.)

A FOAM sliver at foam.gpolab.bbn.com
has been created:

```
Sliver URN: urn:publicid:IDN+ch.geni.net:gpo-infra+slice+gpoN15:ce46ec0e-e297-4569-bb21-  
b7a370e1afbe  
Owner URN: urn:publicid:IDN+ch.geni.net+user+jbs  
Owner e-mail: jbs@bbn.com
```

If the sliver was not automatically approved or rejected, please check
FOAM's pending queue to approve or reject it.

- See the wiki for more info
 - <http://groups.geni.net/geni/wiki/OpenFlow/FOAM/AdminIntro#E-mail>
 - <http://groups.geni.net/geni/wiki/OpenFlow/FOAM#Othersettings>

- Configuration
 - Some guidelines at <http://groups.geni.net/geni/wiki/OpenFlow/FOAM#Initialconfiguration>
 - This should already be done on racks
- Upgrades
 - GPO posts to response-team when it's time to upgrade
 - Rack teams handle upgrades in racks
- FlowVisor
 - <http://groups.geni.net/geni/wiki/FlowVisor#Usefulcommands>
 - You shouldn't need these except when debugging a problem
 - Same upgrade procedures as for FOAM
- When in doubt: gpo-infra@geni.net