

OpenFlow at GEC10

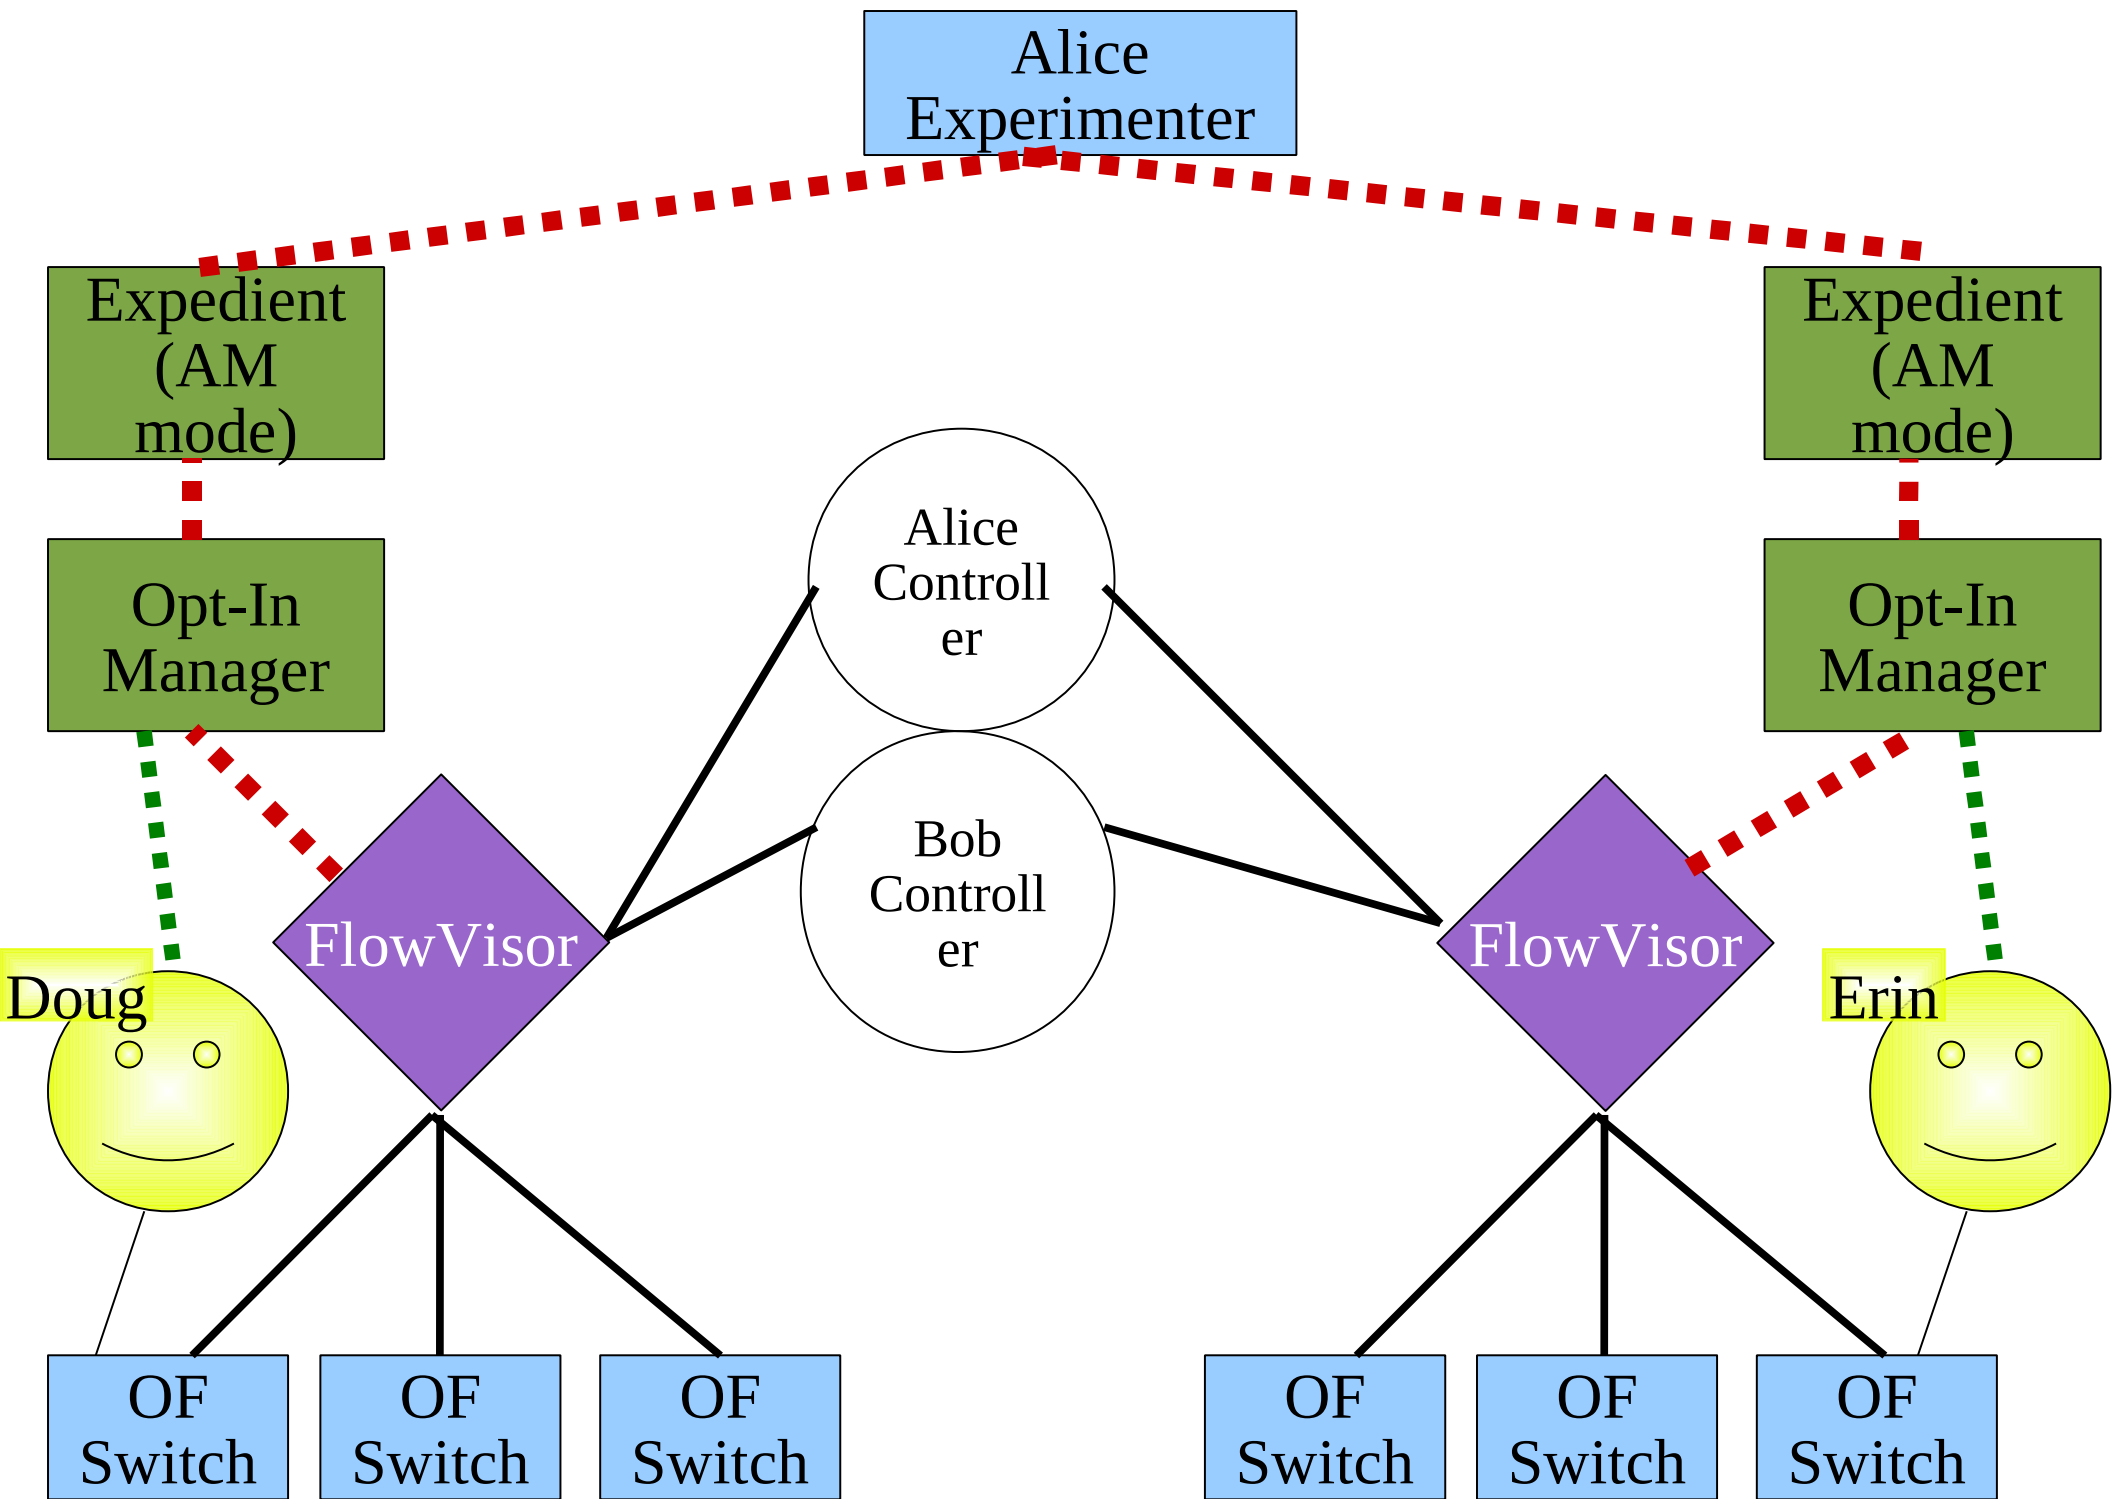
Rob Sherwood

March 16, 2011

Outline

- Overview of OpenFlow on GENI
 - FlowVisor, Opt-In Manager, Expedient
- OpenFlow and Rspecs
 - Current
 - Moving towards ProtoGENI v2 Rspecs
- OpenFlow and Stitching (next talk)
 - Current
 - Comments on proposals

OpenFlow on GENI



Rspecs

OpenFlow Resources

- Switches are nodes
 - 8 byte DPID as unique ID
 - Ports: decimal integers 1-65535
- Links are unidirectional dpid:port pairs
 - $\$sDPID:\$sPort \rightarrow \$dDPID:\$dPort$
- Each link has an allowed “FlowSpace” with ranges
 - $dl_vlan=1235-3000, ip_saddr=128.8.128.0/24, tp_dport=80-80$
- Other info:
 - Controller URL: `tcp://alice.controller.org:6633`
 - User account/passwd for FlowVisor API

Current OpenFlow Rspec

- Advertise
 - No flowspace – Opt-In manager trick
 - Mostly just topology information
 - No bandwidth information
- Request
 - Subset topology (“bound request”)
 - FlowSpace per link: lots of duplication
 - No bandwidth reservations
- Translates to omni_spec (by hand)

Thoughts on ProtoGENI v2 Rspec

- Stick with “bound” requests for now
 - No need for ticket, manifest rspecs?
- Create a FlowSpace extension
 - 12+ OpenFlow tuples: ranges?
 - For links, switches, slices?
- Create a bandwidth reservation extension
- Bottom line:
 - “have rough consensus; now need running code”

Slice Stitching

Current Stitching

- Hard code touch points
 - PlanetLab, ProtoGENI, other OpenFlow aggregates
 - Advertised in Expedient (in Clearing House mode)
- Experimenter “hand” stitches
 - Via rspec using “tree model”
 - Intra-OpenFlow: automagic topology discovery
- Problems
 - Doesn't scale
 - Error prone
 - “Sub-optimal user experience”

OpenFlow and Proposal

- Agree at high-level
- Two level discovery hierarchy
 - Global, inter-AM view
 - Static information: touch points
 - Local, intra-AM view
 - Dynamic information: link status, bandwidth availability
- Likely need to add OpenFlow linktype
 - Assumes VLAN = Bandwidth reservation
 - Describe virtual-wire-like functionality
 - “No horse in this race”

Perceived Hard Problems

- Mixed chain and tree stitching
 - Can we have a common AM-level API?
 - Can a chain node pretend to be the stitching service for a tree node? Security?
- How static is our static data?
 - Can/should it be dynamically inferred?
 - How do we detect errors?
- How similar is this to an IGP/EGP split?
 - Assumes intra-AM connectivity

Comments on Rob Ricci's Slides

- Two fundamental use cases
- Loose source routing
 - Hide technical details
 - Experimenter: Just make it happen
 - “Chain mode”
- Strict source routing
 - Expose technical details
 - Experiment: I care about the topology
 - “Tree mode”
- Use cases motivate different solutions

Comments on Jeff's Slides

- DAG might be a solution to tree + chain mixing
 - Inherently?
- This problem still suffers from conflicting naming conventions
 - Not directed at Jeff in particular
 - Largely historical
 - Potential point for GPO leadership

Conclusions

- Rspec and stitching proposals seem like they should work
 - read: “have rough consensus”
- Proof is in pudding
 - read: “need running code”
- Flexibility of extension mechanisms critical