Enterprise Geni Update

Rob Sherwood, Guido Appenzeller
Stanford Clean Slate Lab

Chicago, June 2009

Ethernet Switch

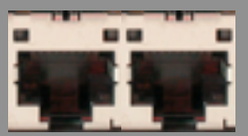

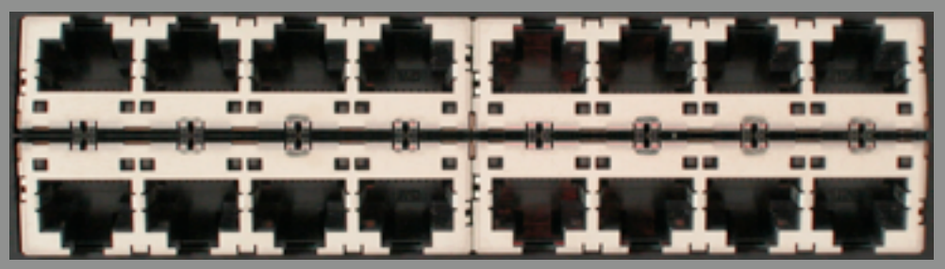

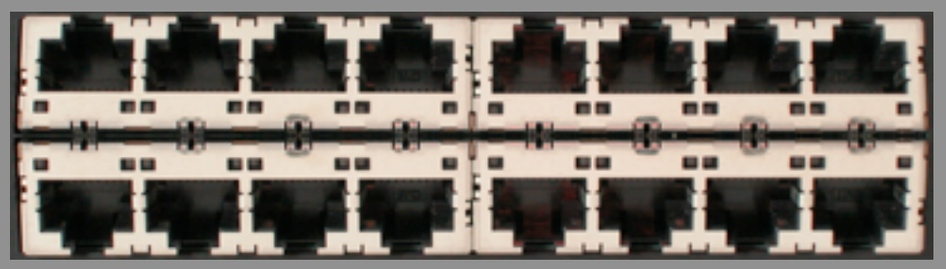

Control Path (Software)

Data Path (Hardware)

Control Path

Data Path (Hardware)

Control Path

Data Path (Hardware)

Control Path

Data Path (Hardware)

Control Path

OpenFlow

Data Path (Hardware)

OpenFlow Controller

Control Path

OpenFlow

Data Path (Hardware)

OpenFlow Controller

OpenFlow Protocol (SSL)

Control Path

OpenFlow

Data Path (Hardware)

FLOWVISOR BASED VIRTUALIZATION

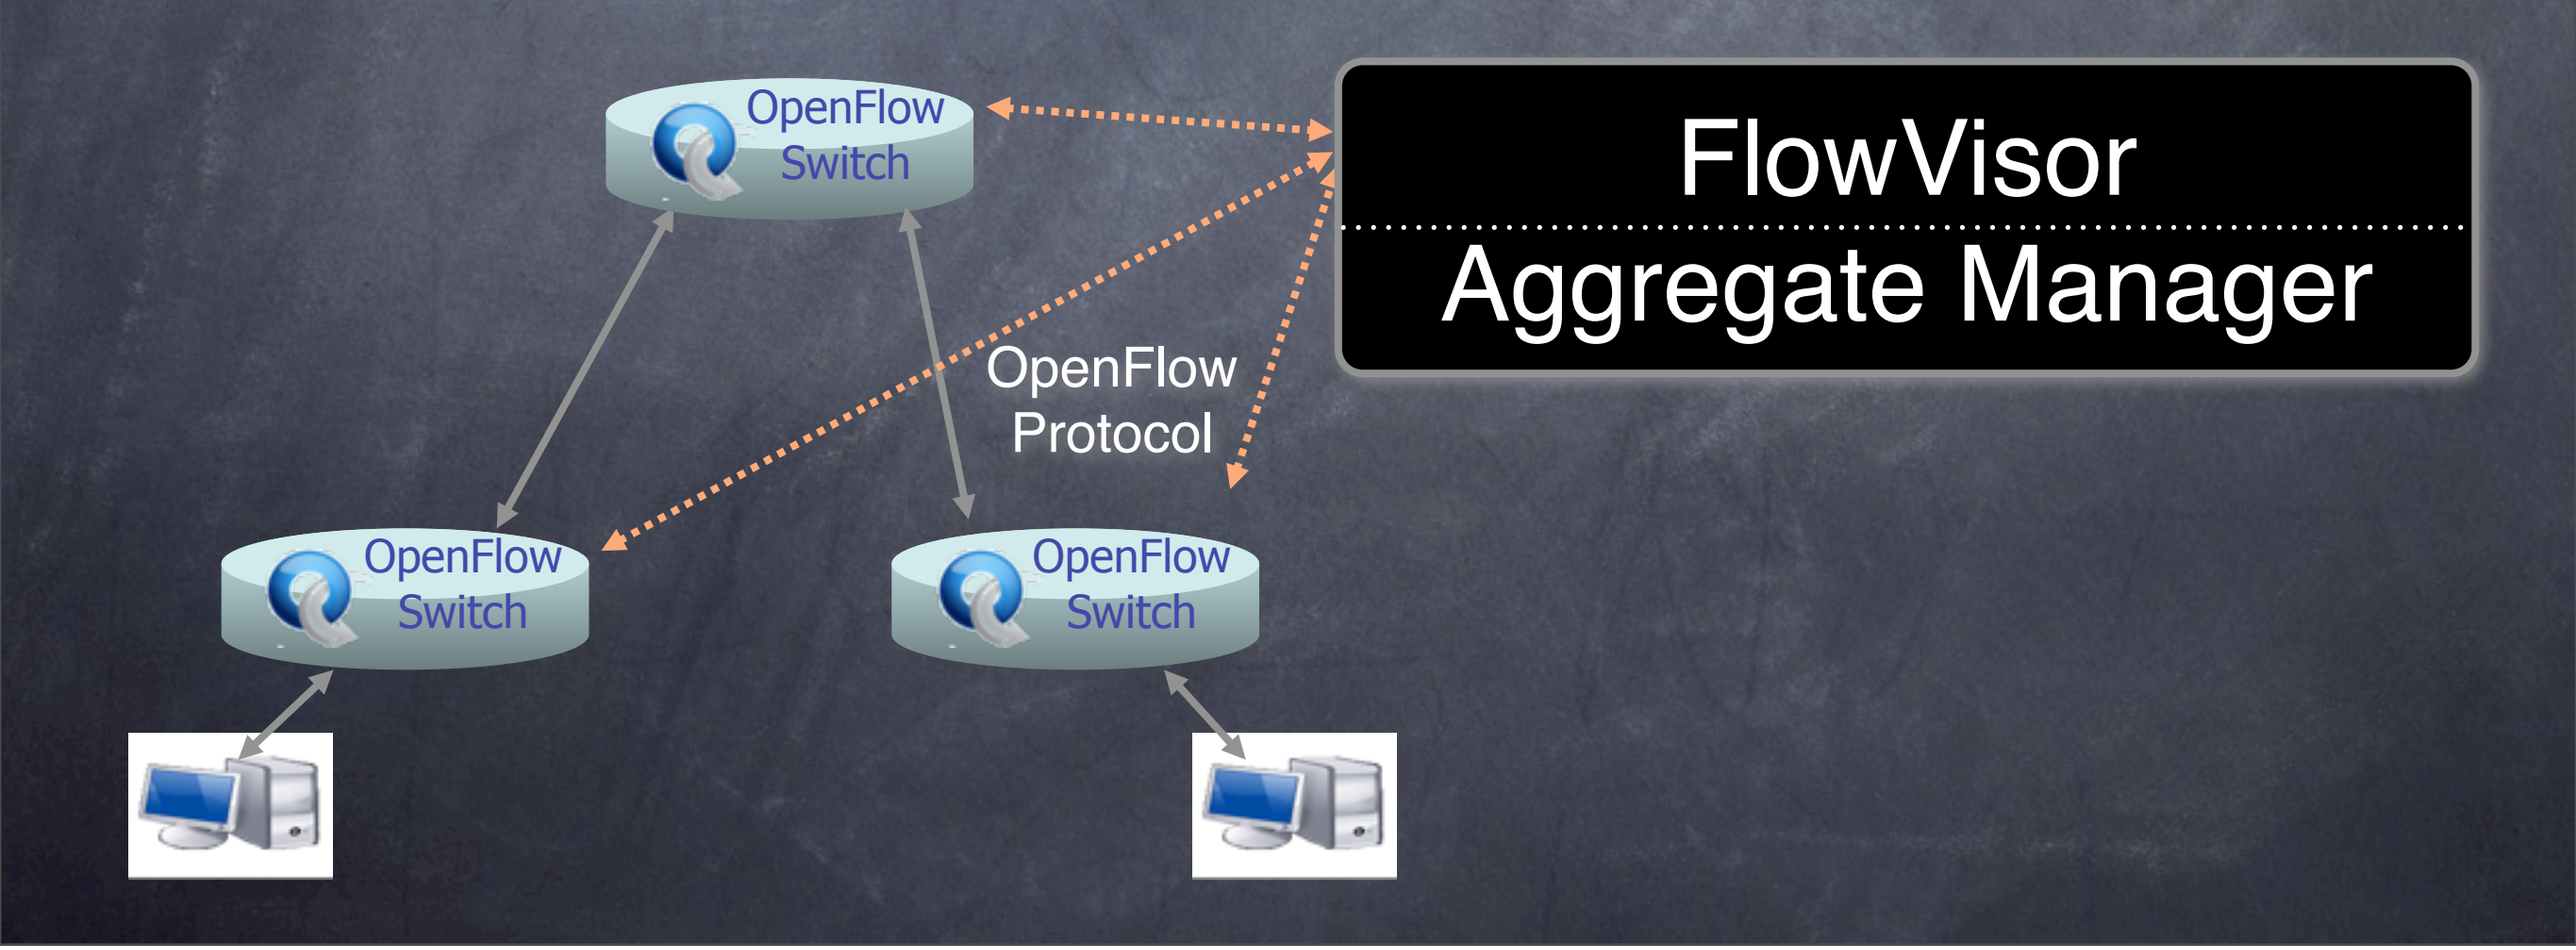

FLOWVISOR BASED VIRTUALIZATION

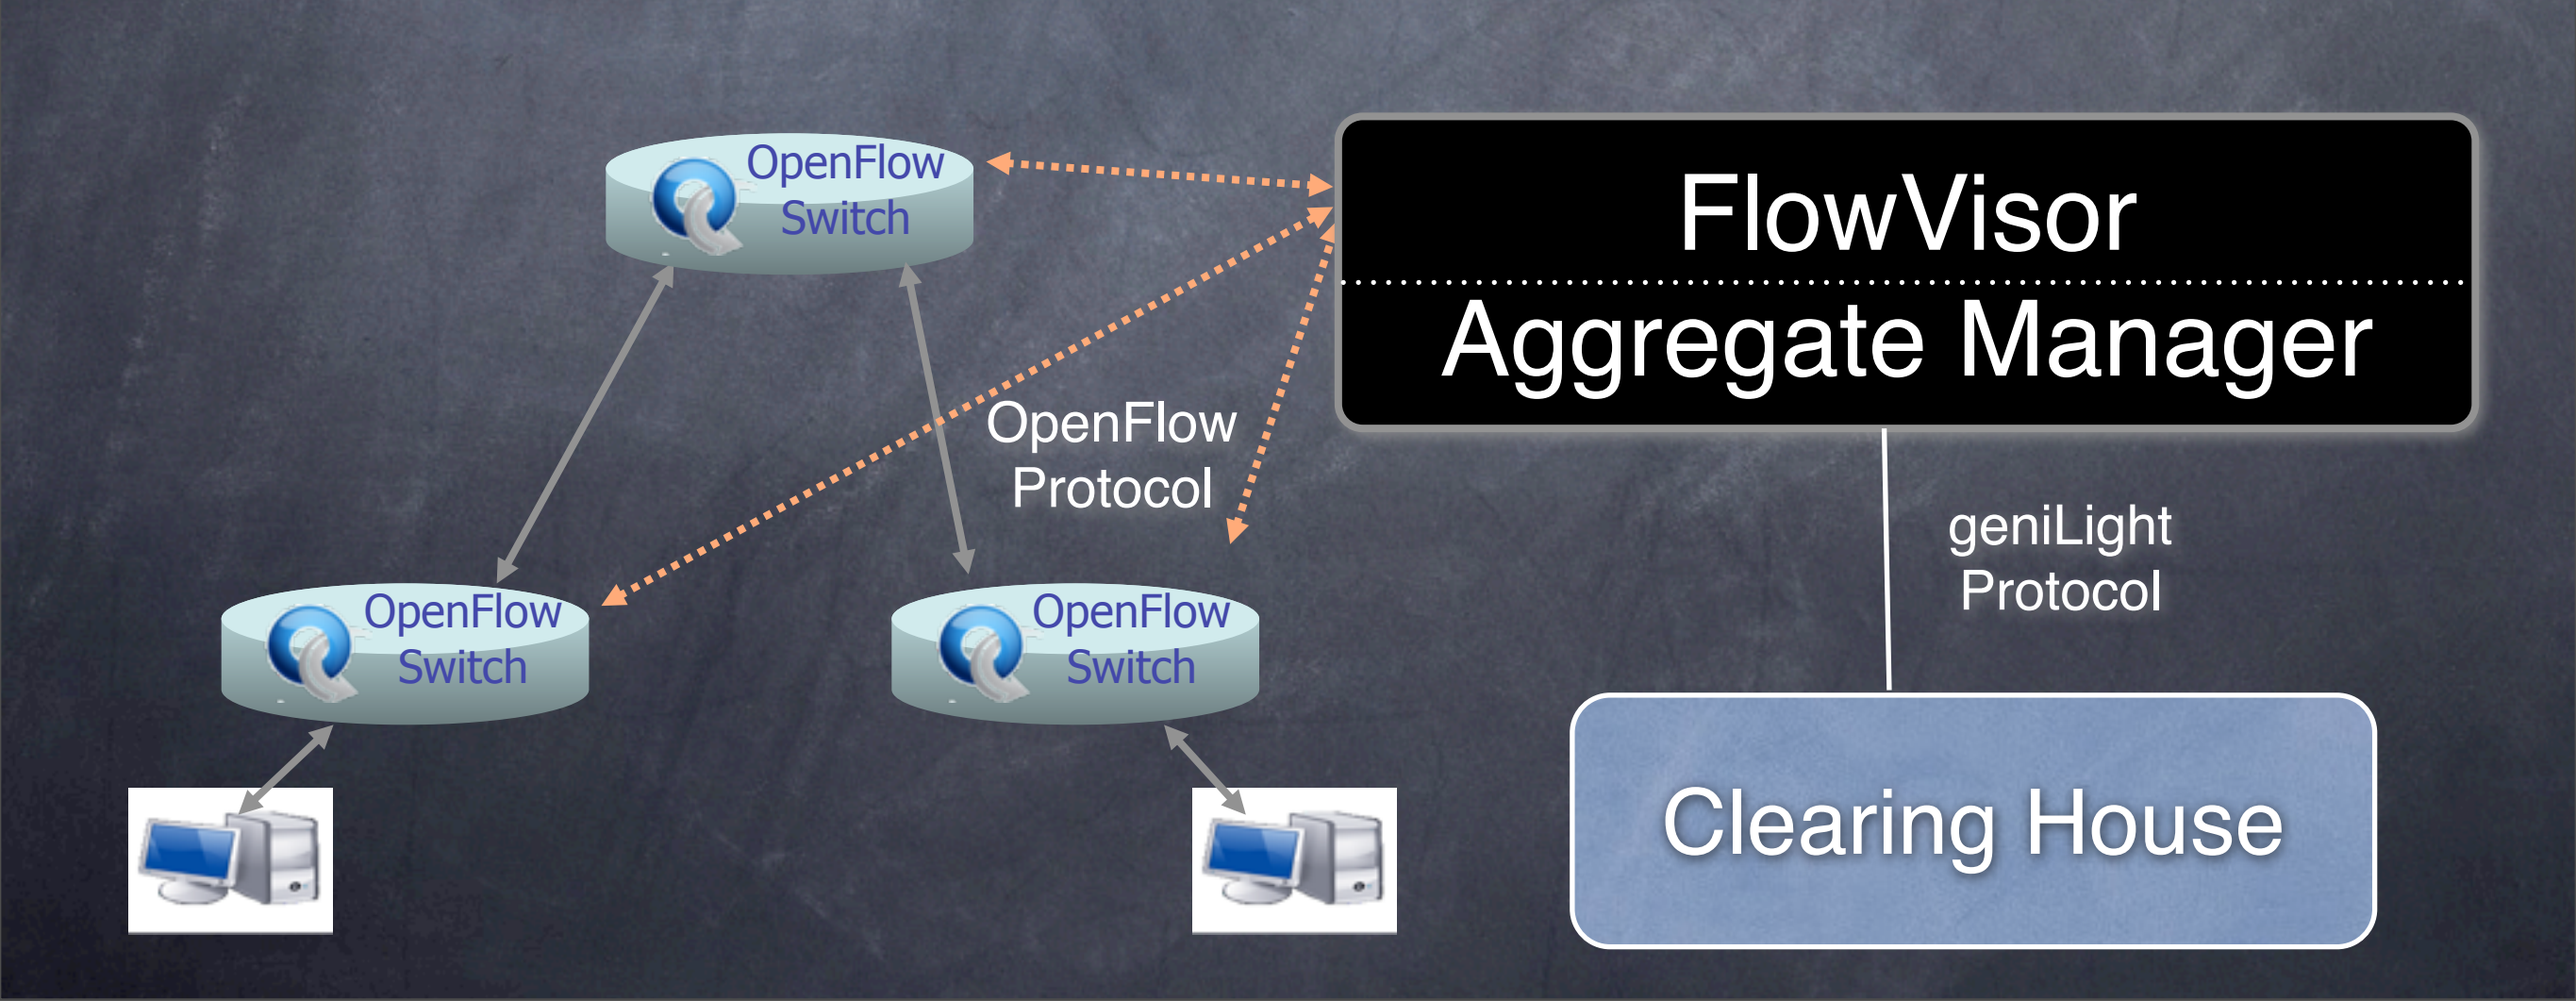

FLOWVISOR BASED VIRTUALIZATION

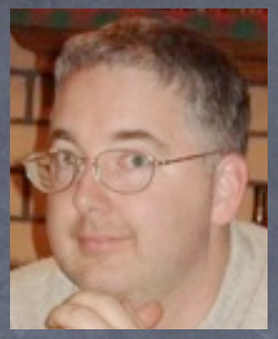

Aaron's Controller

OpenFlow Protocol

OpenFlow Switch

OpenFlow Switch

OpenFlow Switch

OpenFlow Protocol

FlowVisor

Aggregate Manager

geniLight Protocol

Clearing House

FLOWVISOR BASED VIRTUALIZATION

Aaron's Controller

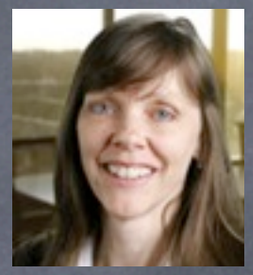

Heidi's Controller

Craig's Controller

OpenFlow Protocol

OpenFlow Switch

OpenFlow Switch

OpenFlow Switch

OpenFlow Protocol

FlowVisor

Aggregate Manager

geniLight Protocol

Clearing House

Current Status

- FlowVisor: mostly feature complete
  - publicly released: June 2nd
- Aggregate Manager:
  - Resource discovery; reports to CH as rspec
  - Accepts reservations; converts rspec to Flowvisor config
- CH: Implemented toy clearing house for testing: integrate later

Current Status: Rspec

- Big Picture: Work in Progress
- Rspec requirements still in flux
- Integration likely premature

Rest of Talk: E-Geni Rspec

- Switches
- Interfaces
- “FlowSpace”
- Opt-In
- Inter-Aggregate connectivity

Rspec: Switches

- Switches:
  - Identify via “datapath ID” (i.e., MAC addr)
  - List of interfaces
  - Unlike a node, we don’t login to switches

Rspec: Interfaces

- Interfaces: think unidirectional links
  - List of remote interfaces for connectivity
    - i.e., topology information
  - FlowSpace policy per interface
  - Policies are enforced on egress

Rspec: FlowSpace

- FlowSpace: header field=value pairs + action
- Packet classifier: part of slice definition
- Header fields: ip_src, ip_dst, ethertype, etc.
- Actions: allow, deny, listen-only
- Example: all web traffic except to main server
  - ip_src=1.2.3.4 tcp_dport=80 :: DENY
  - ip_src=1.2.3/24 tcp_dport=80 :: ALLOW

E-Geni Opt-In: Use Case

- Alice creates an experimental slice
  - Alice specifies class of traffic for Opt-In
  - Alice sends Bob an email to Opt-In
  - Bob follows URL, reads terms, clicks OK
  - Corresponding class of Bob's traffic is now controlled by Alice

Rspec: Opt-In

- Express what users experimenter wants?
  - “All”, “First 10”, “only port 7 on switch 3”
- Description of Experiment
  - Privacy, SLA, etc.
  - Class of traffic: i.e., FlowSpace
  - URL?

Inter-Sliver Compatibility

- Use case: “Give me four PlanetLab nodes and the E-Geni network that connects them”
- Need to know PL network points of attachment
- Need to communicate user’s e2e network slice
  - On PL: bind() dynamically maps L4 ports to an experimenter’s sliver
  - Network slice should be consistent e2e
- How do we propagate this info? Rspec?

Conclusion

- We have a fair amount of running code
  - But have a lot more to go...rspec still in flux
- We're currently working on Opt-In
- Inter-sliver compatibility needs some thought

Beware: Backup Slides!

(Old) Talk Outline

- Status Update
- Current challenges:
  - Opt-In
- Future challenges:
  - Inter-aggregate information exchange

E-Geni Slice =~ Topology + FlowSpace

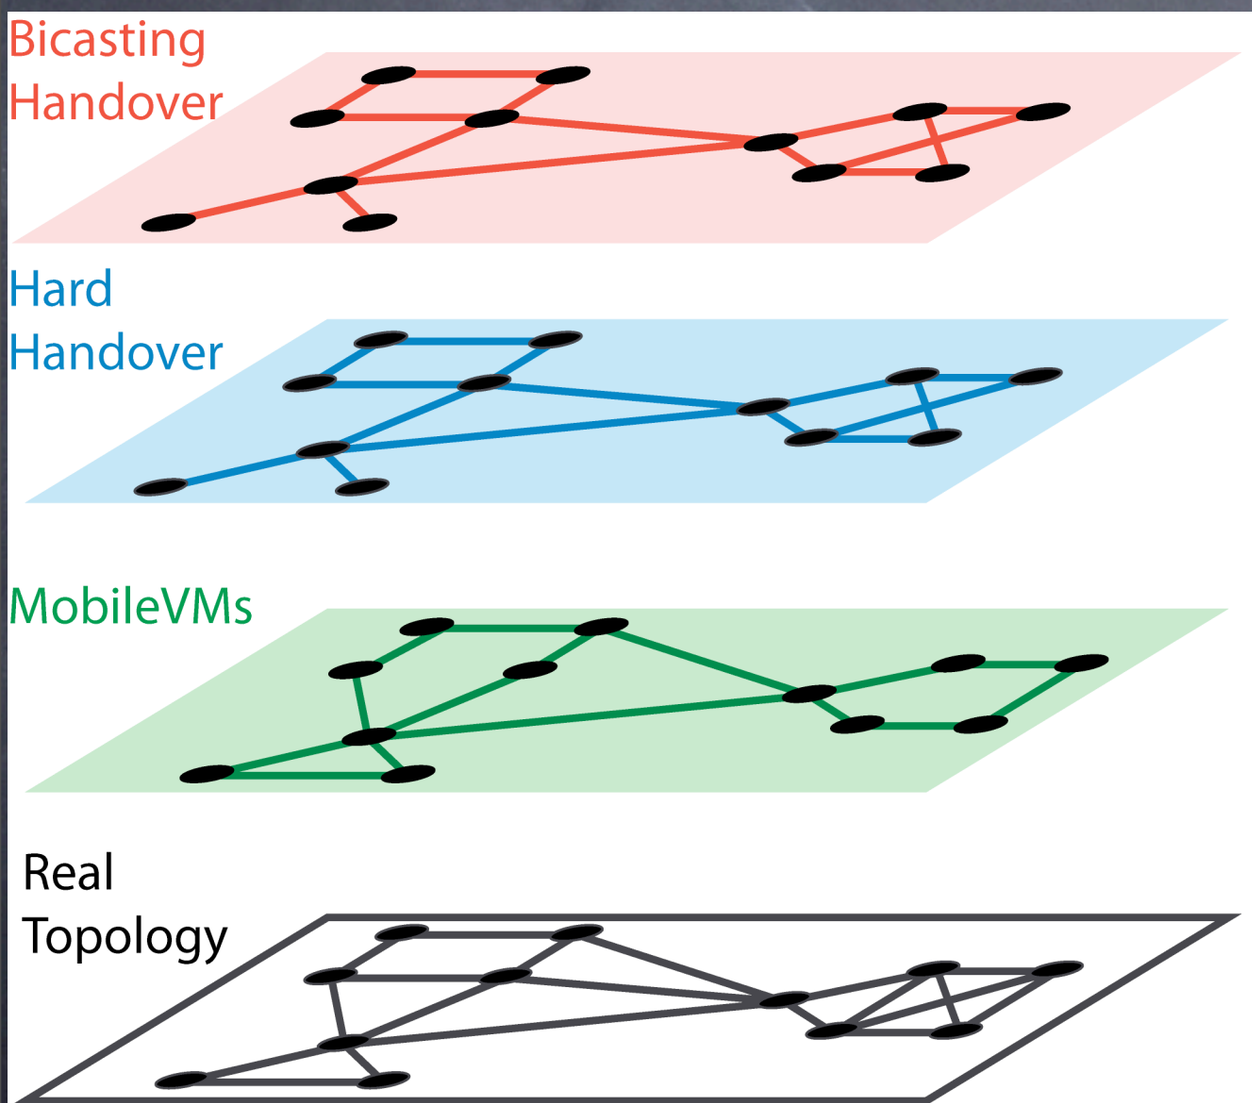

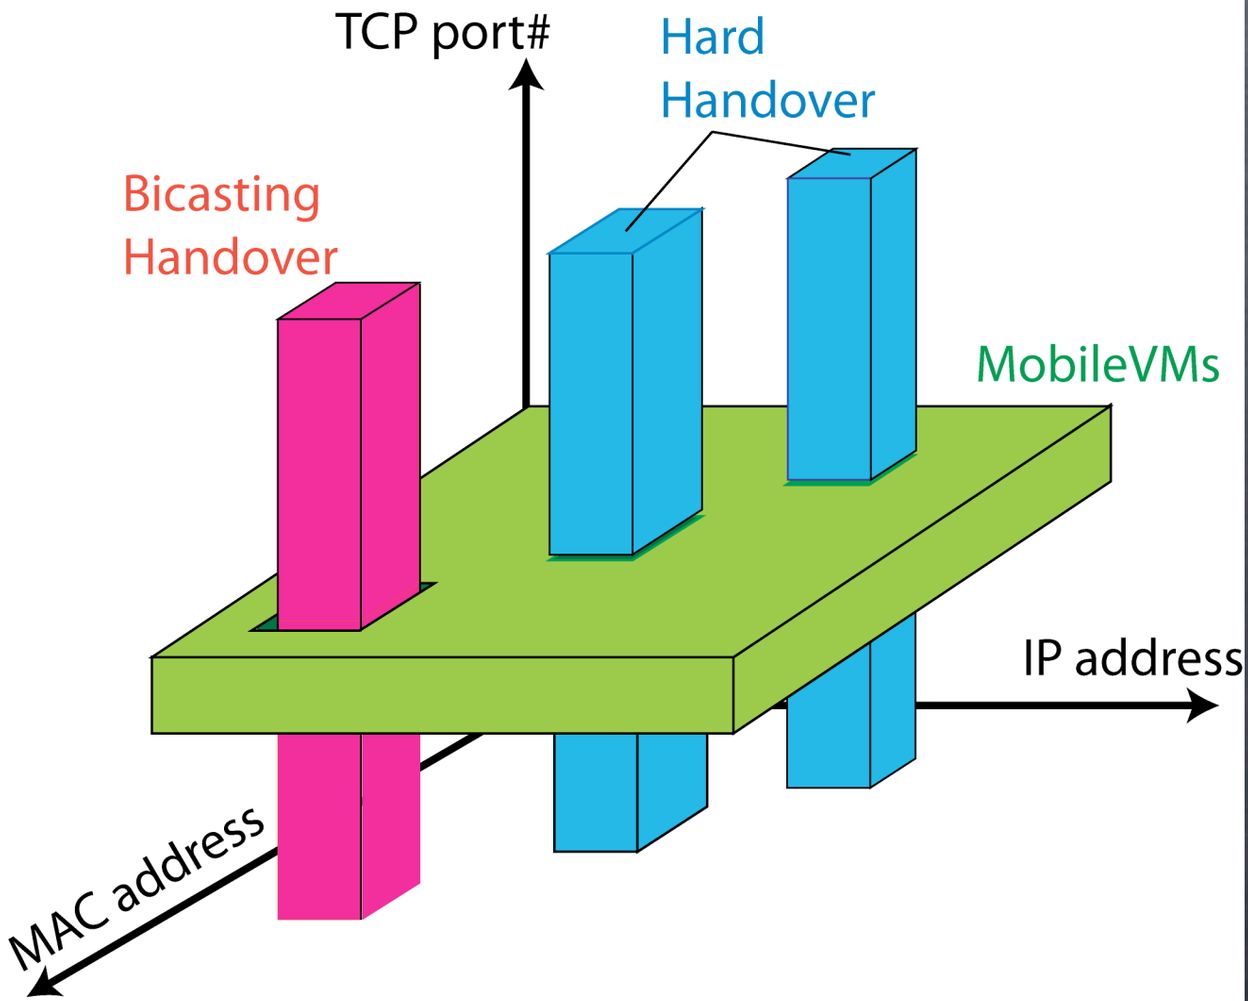

Opt-In Challenges (1/2)

How do we map packets to slices?

- Mark packets explicitly
  - VLAN tags? TOS bits? IP Subnet?
  - Hard end-to-end
  - Doesn’t work for all experiments
- Keep explicit mapping of flows to experiments
  - Hard to aggregate/scale

Opt-In Challenges (2/2)

- Verify Bob has permission to delegate traffic
  - Authentication policy? To whom?
- Stop Bob from Opting In Cathy
  - Accidentally: dhcp turn over
  - Maliciously: trivial packet sniffer
- Don’t block Opt-Out channel!
- Move users from production to OpenFlow VLAN

E-Geni and Other Aggregates?

- Need two mechanisms:
  - Learn information:
    - Static Config? Discovery Protocol?
  - Propagate to CH/experimenter
    - Rspec? Something else?

E-Geni Architecture

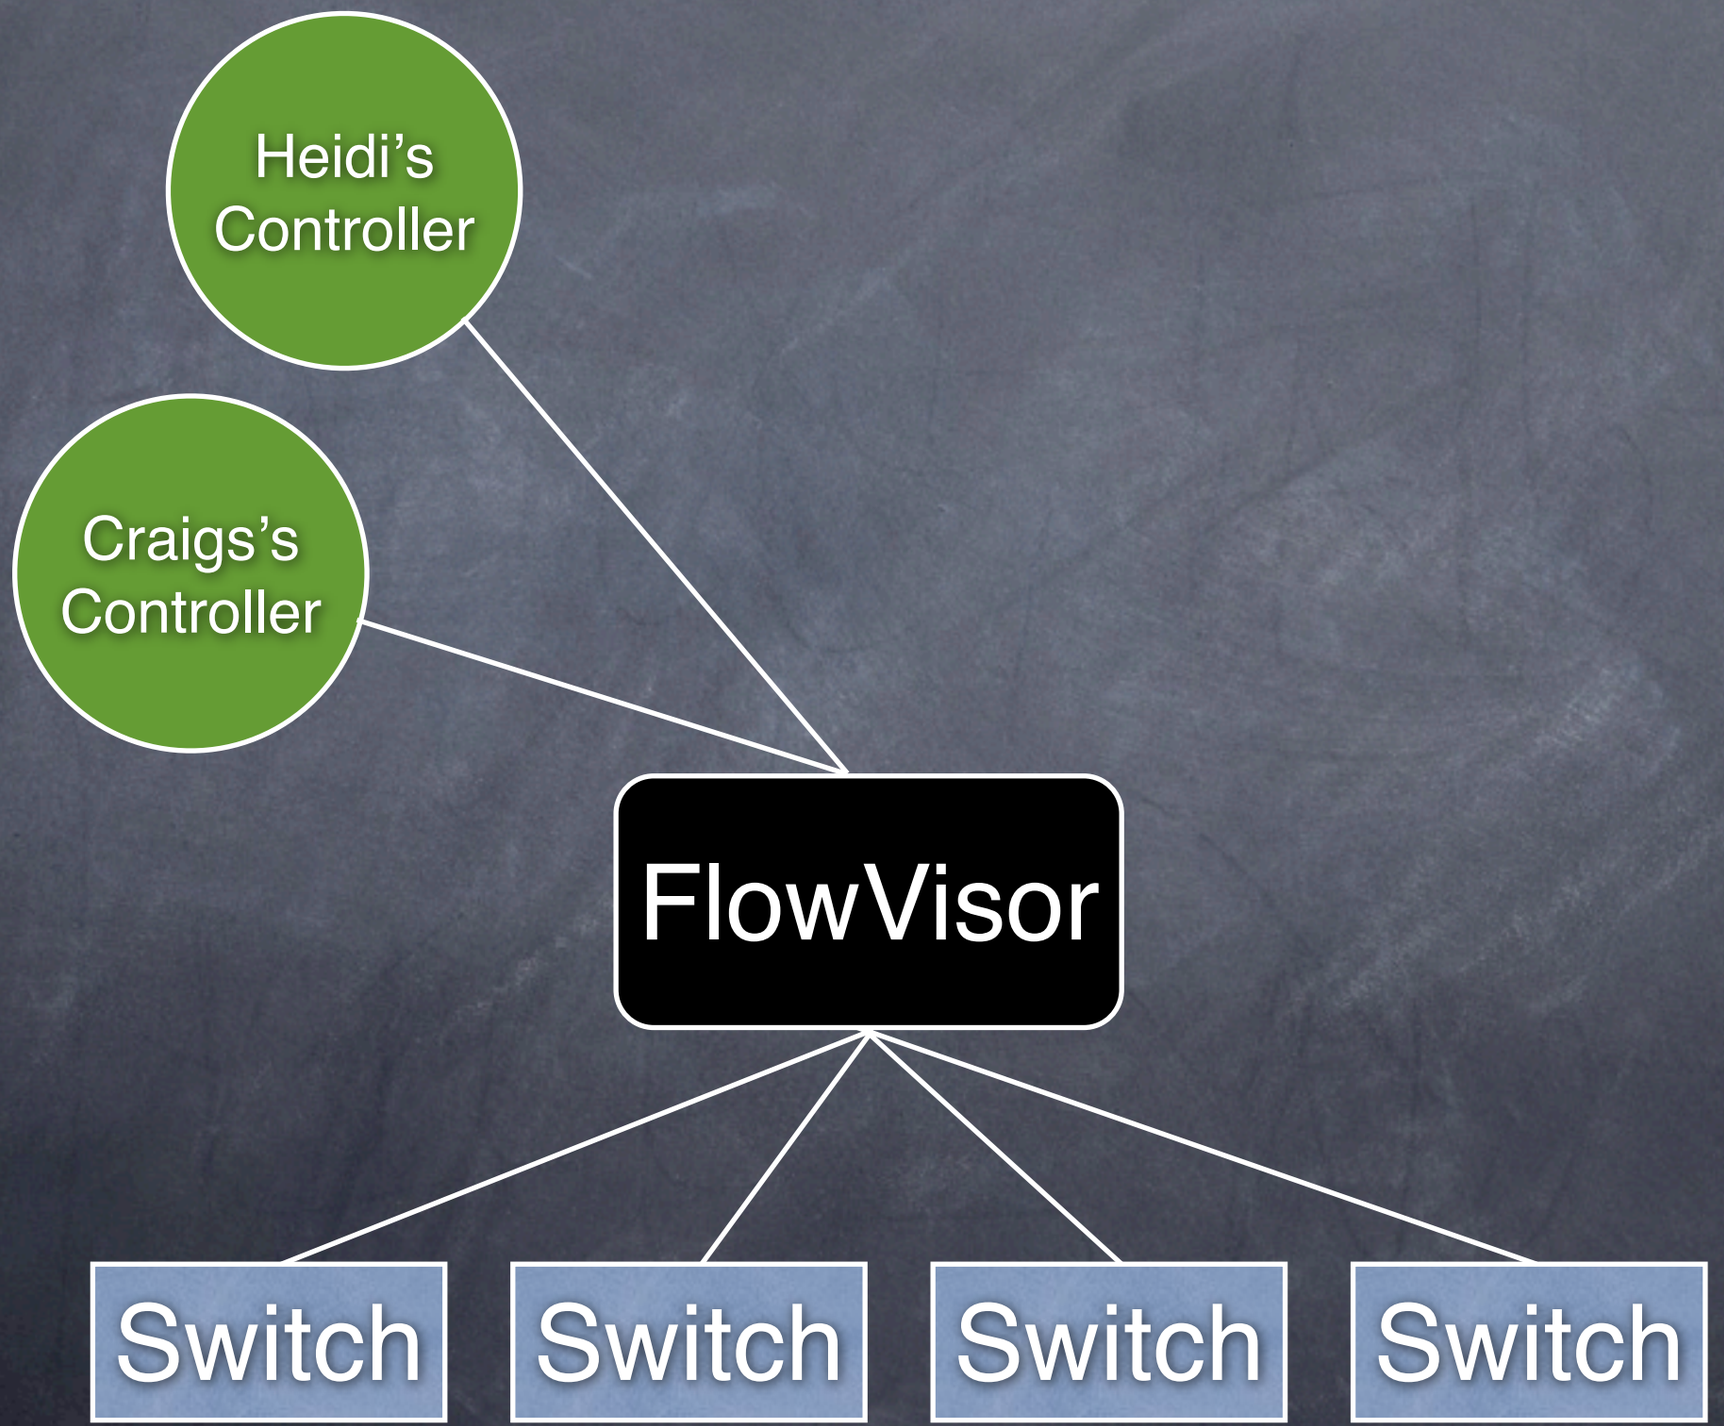

E-Geni Architecture

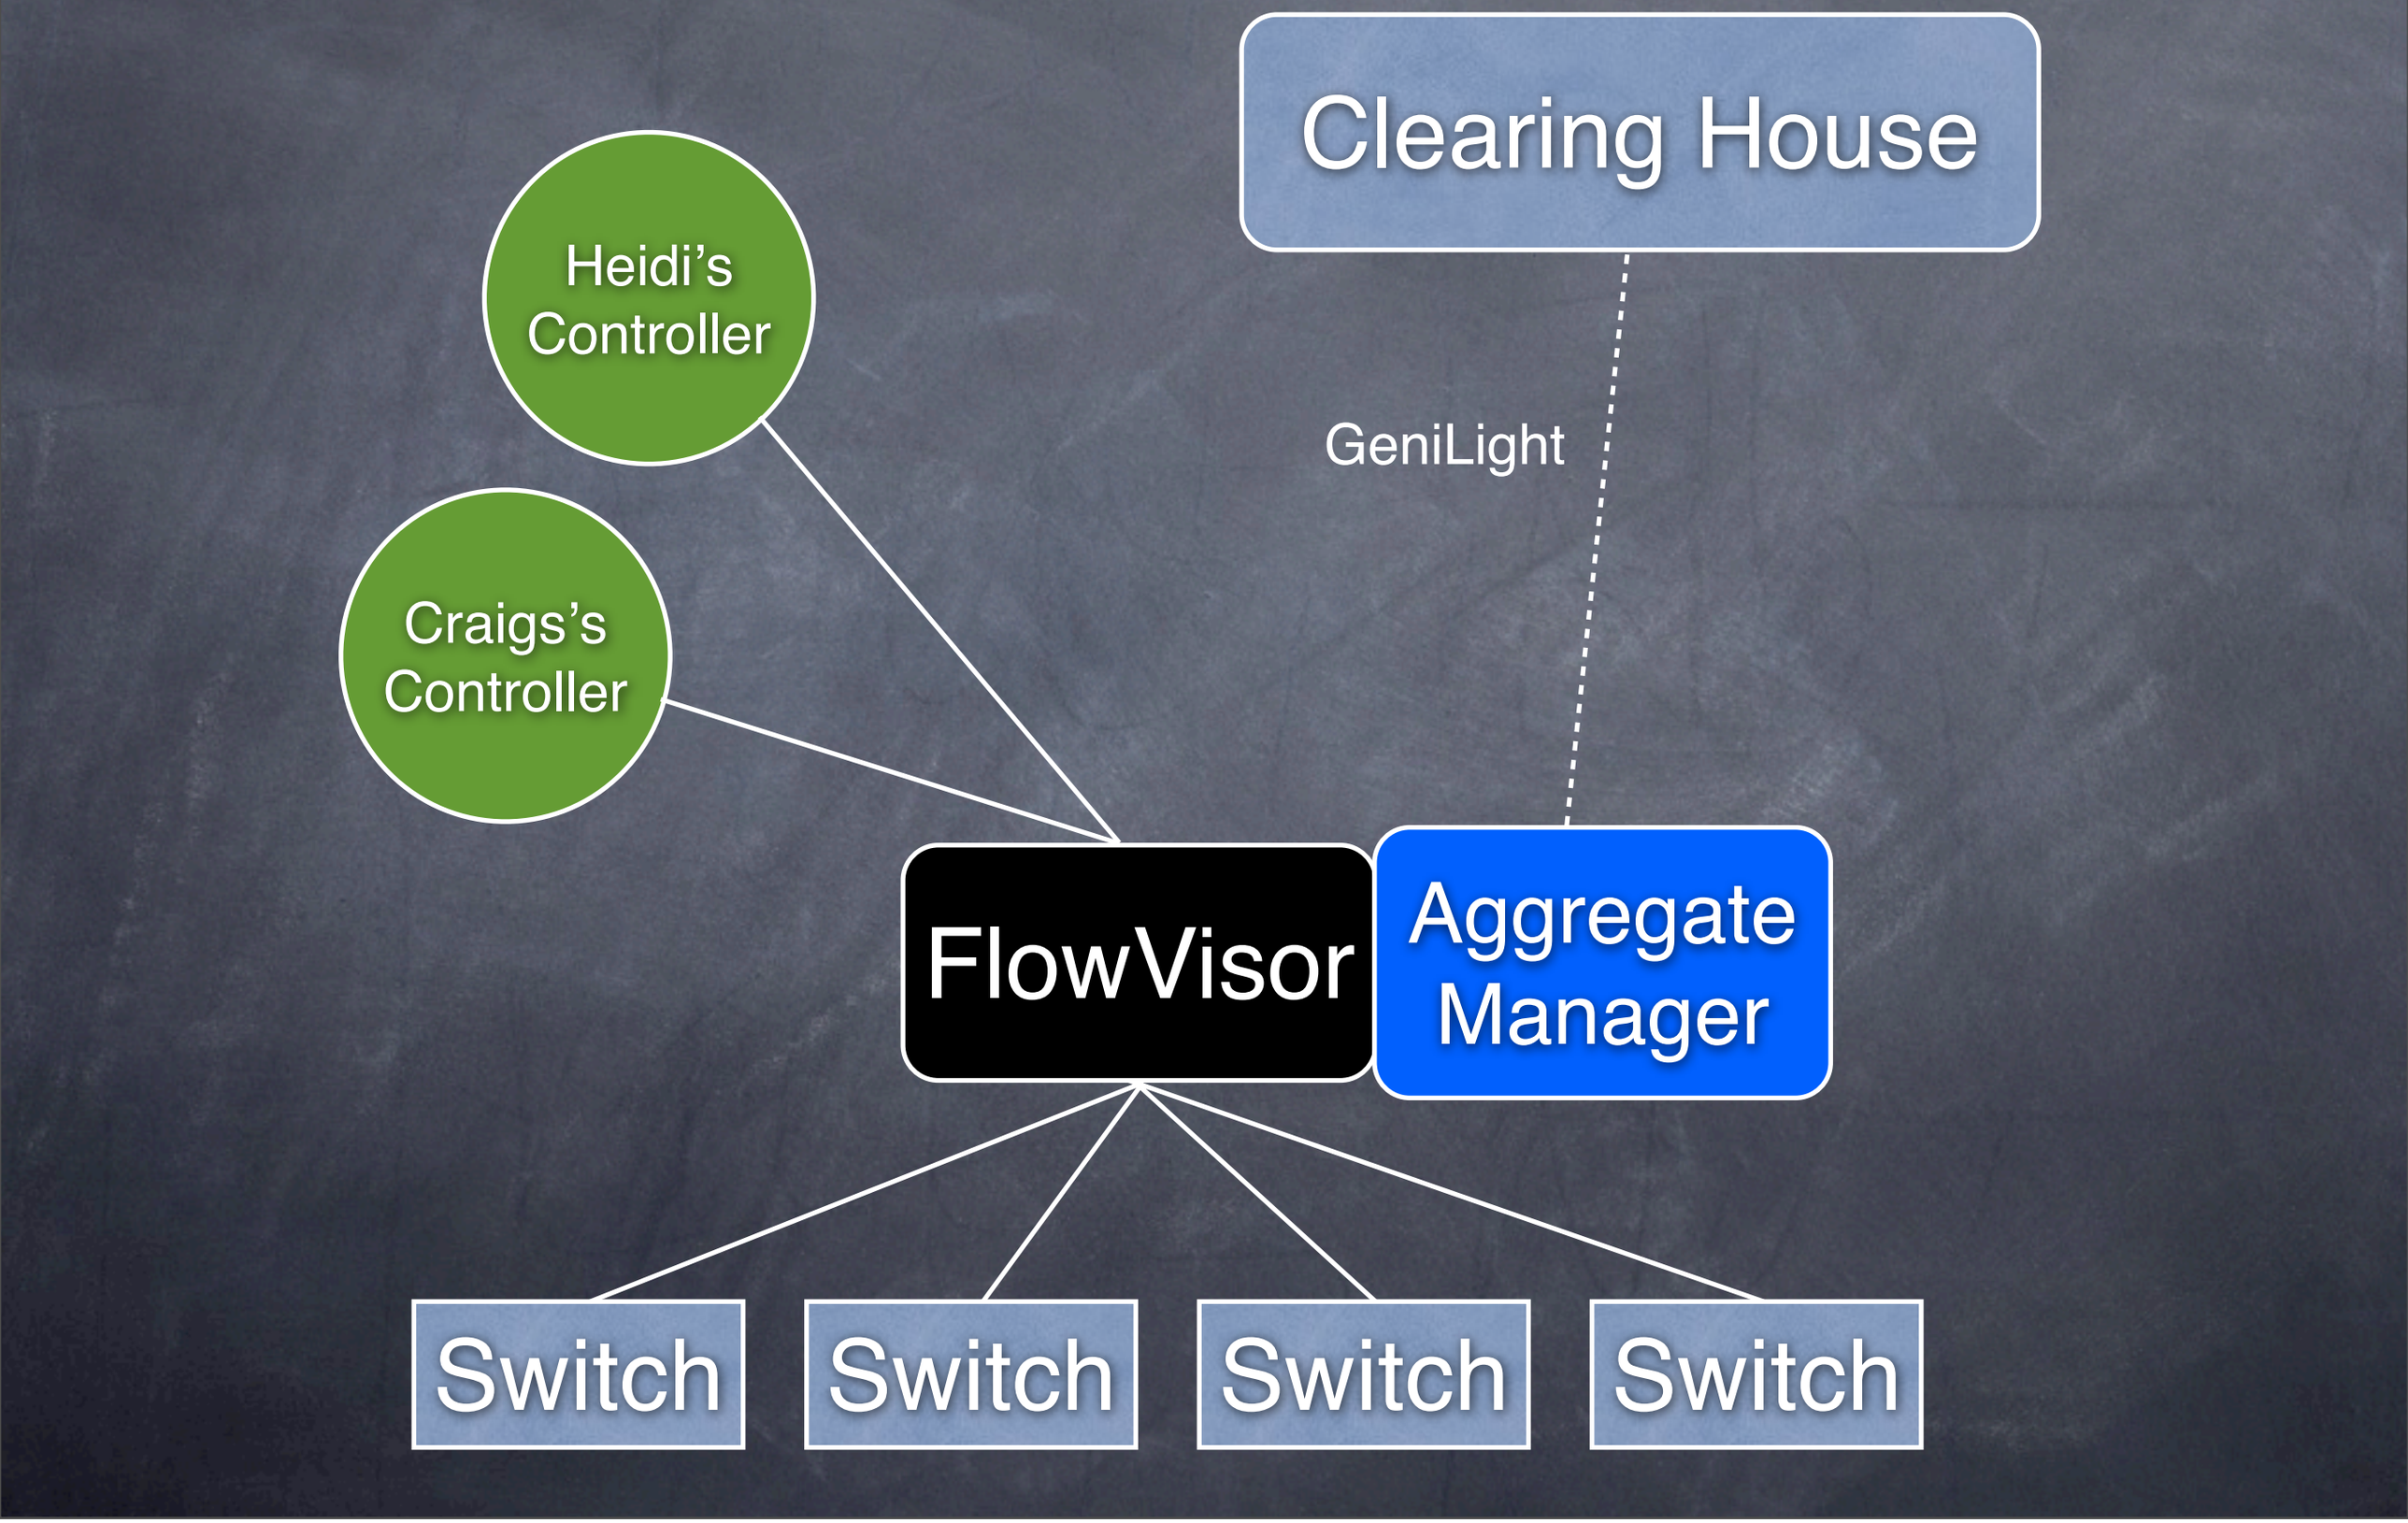

E-Geni Architecture