
Perspectives on the GENI Control Plane

Jeff Chase


GEC 2 NSF 3/4/08



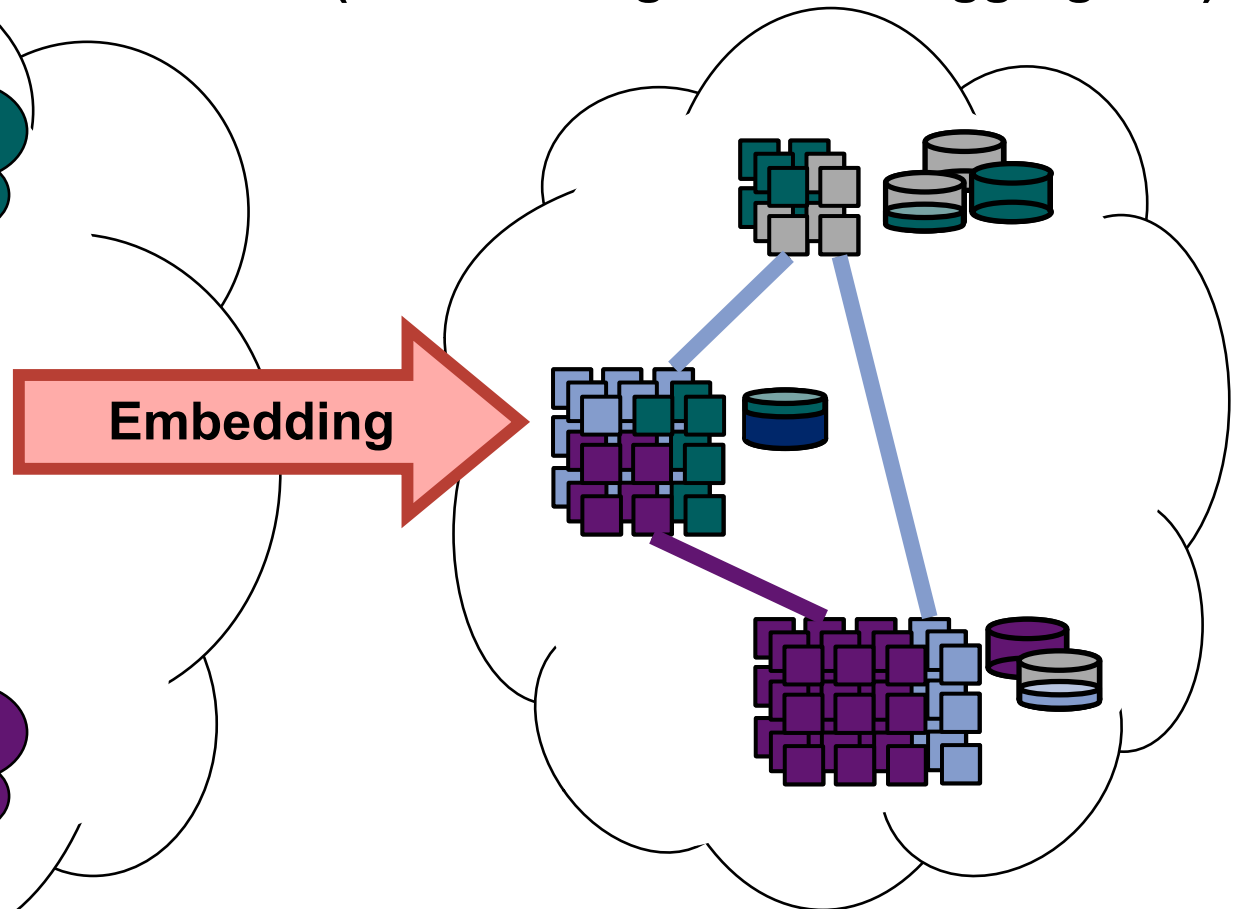
D u k e S y s t e m s

Setting the Stage

Experiments (Guests)



Sliverable GENI Substrate (Contributing domains/Aggregates)



Define: Control Plane

GGF+GLIF: "Infrastructure and distributed intelligence that controls the establishment and maintenance of **connections** in the network, including protocols and mechanisms to disseminate this information; and algorithms for automatic delivery and on-demand provisioning of an **optimal path between end points.**"

s/connections/slices/

s/optimal path/embedded slices

provisioning += and programmed instantiation



Global Lambda Integrated Facility



Define: Control Plane

Control plane defines “the set of entities that interact to establish, maintain, and release resources and provide...connection control functions”.


At bottom this is about modularity and interfaces.

If the entities are owned, operated, and developed by different stakeholders over long time spans, then we call it “architecture”.



The GENI Control Plane

- Programmable sliver allocation + control
 - Delegation of authority etc.
 - Instrumentation (feedback)
- Resource representation and exchange
- Foundation for resource discovery, path discovery, topology discovery



Design Tensions

- Governance vs. freedom
- Coordination vs. autonomy
- Control vs. manageability
- Assurance vs. robustness
- Predictability vs. efficiency
- Quick vs. right
- Inclusion vs. entanglement
- Etc. etc. ...



Design Tensions, continued

- *What is standardized vs. what is open to innovation?*
- I am taking an aggressive view about openness to innovation.
 - Because we want it to last a long time
 - Because that is what it is for
- Standardization kills innovation.
 - “Apocalypse of the two elephants” [Clark]



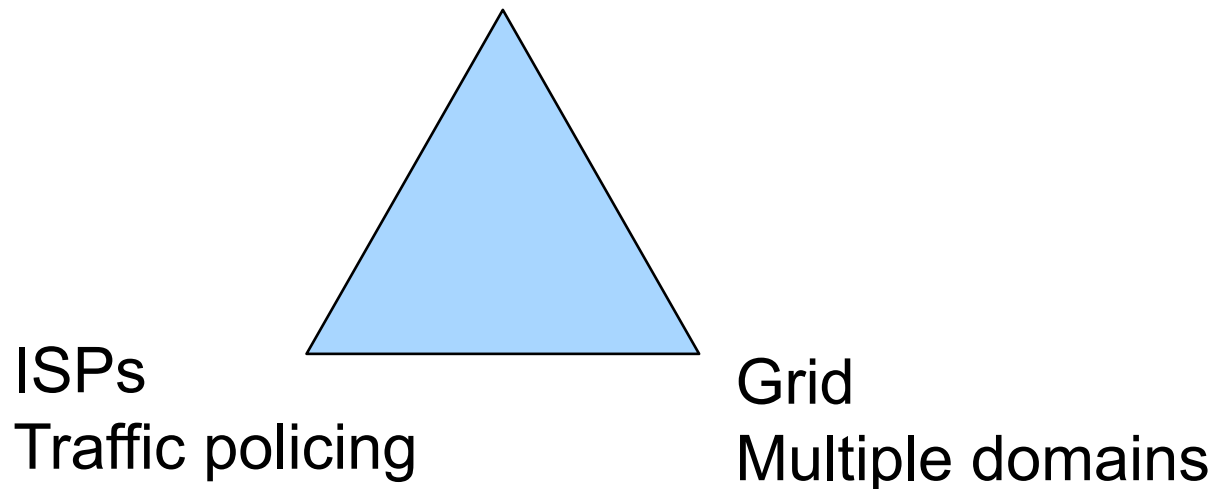
Modularize Innovation

- Control plane design should enable local innovation within each entity.
- Can GENI be a platform for innovation of platforms? Management services?
 - Enhance PlanetLab's unbundled management
- A plea for pluralism...
- Key: no one-size-fits-all standard for information exchange.



A Platform for “Utility Networking”

Hosting centers (utilities)
Untrusted/untrusting guests
Virtualization and isolation



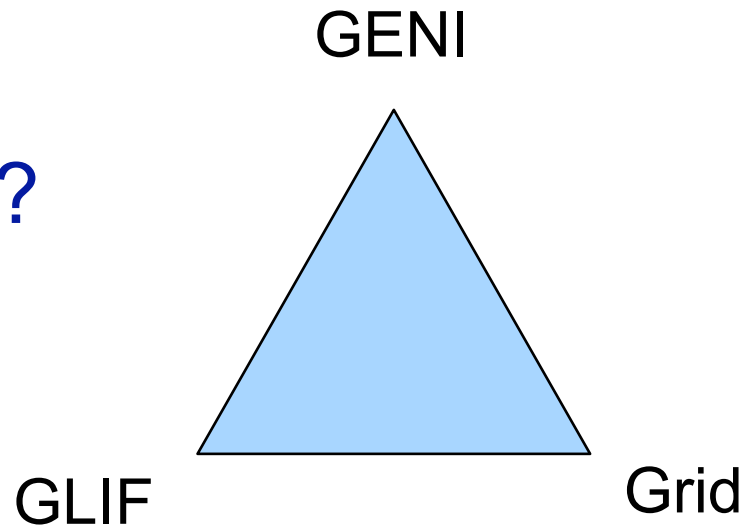
Experiences Building PlanetLab, OSDI 06



A Platform for “Utility Networking”

Programmable substrate
Dynamic end-to-end slivering

Cooperate
or compete?



Global Lambda Integrated Facility

OASIS  **DCML** Standardizing Information Exchange
for the Data Center Environment



Network Description Language

```
<ndl:Interface rdf:about="#tdm3.amsterdam1.netherlight.net:501/3">
  <ndl:name>tdm3.amsterdam1.netherlight.net:501/3</ndl:name>
  <ndl:connectedTo
    rdf:resource="http://networks.internet2.edu/manlan/manlan.rdf#manlan:if1"/>
  <ndl:capacity
    rdf:datatype="http://www.w3.org/2001/XMLSchema#float">1.244E+9</ndl:capacity>
</ndl:Interface>
<ndl:Interface rdf:about="http://networks.internet2.edu/manlan/manlan.rdf#manlan:if1">
  <rdfs:seeAlso rdf:resource="http://networks.internet2.edu/manlan/manlan.rdf"/>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:ndl="http://www.science.uva.nl/research/sne/ndl#"
  xmlns:geo="http://www.w3.org/2003/01/geo/wgs84_pos#">
  <!-- Description of Netherlight -->
  <ndl:Location rdf:about="#Amsterdam1.netherlight.net">
    <ndl:name>Netherlight Optical Exchange</ndl:name>
    <geo:lat>52.3561</geo:lat>
    <geo:long>4.9527</geo:long>
  </ndl:Location>
  <!-- TDM3.amsterdam1.netherlight.net -->
  <ndl:Device rdf:about="#tdm3.amsterdam1.netherlight.net">
    <ndl:name>tdm3.amsterdam1.netherlight.net</ndl:name>
    <ndl:locatedAt rdf:resource="#Amsterdam1.netherlight.net"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/1"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/2"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/3"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:501/4"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:502/1"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:502/2"/>
    <ndl:hasInterface rdf:resource="#tdm3.amsterdam1.netherlight.net:502/3"/>
```



Key Questions

- Who are the entities (actors)?
- What are their roles and powers?
- Whom do they represent?
- Who says what to whom?


Control plane defines “the set of entities that interact to establish, maintain, and release resources and provide...connection control functions”.



GENI Principals



Researcher: A user that wishes to run an experiment or service in a slice, or a developer that provides a service used by other researchers.



A **slice authority (SA)** is responsible for the behavior of a set of slices, vouching for the users running experiments in each slice and taking appropriate action should the slice misbehave.



A **management authority (MA)** is responsible for some subset of substrate components: providing operational stability for those components, ensuring the components behave according to acceptable use policies, and executing the resource allocation wishes of the component owner.




GENI Principals



Researcher: A user that wishes to run an experiment or service in a slice, or a developer that provides a service used by other researchers.

Programmatic?
Automated?



A **slice authority (SA)** is responsible for the behavior of a set of slices, vouching for the users running experiments in each slice and taking appropriate action should the slice misbehave.

Who does it represent?
Policing or helper?



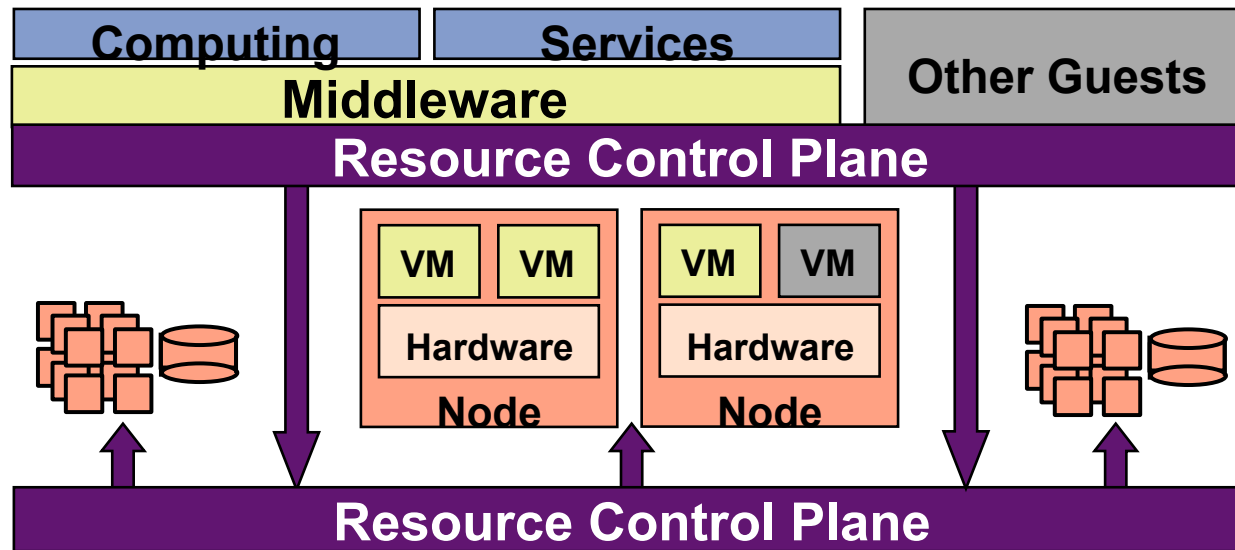
A **management authority (MA)** is responsible for some subset of substrate components: providing operational stability for those components, ensuring the components behave according to acceptable use policies, and executing the resource allocation wishes of the component owner.

Who does it represent?
Policing or helper?

Aggregates? Clearinghouse? GSC?



Open Resource Control Architecture (Orca)



- Contract model for resource management
- Programmatic interfaces and protocols
- Automated lease-based allocation and assignment
- Share infrastructure among multiple guests
- <http://nicl.cod.cs.duke.edu/orca/>



Who Are the Actors?

- Principle #1: Entities (actors) in the architecture represent the primary stakeholders.
 1. Resource **owners**
 2. Slice owners/controllers (**guests**)
 3. The **facility** itself, or resource scheduling services acting on its behalf.
- Others (e.g., **institutions**) are primarily endorsing entities in the trust chains.



Contracts

- Principle #2: provide pathways for **contracts** among actors.
 - **Accountability**
- Be open with respect to what promises an actor is *permitted* to make.
 - **Open innovation for contract languages and tools**
 - **Yes, need a least one least common denominator**
 - RSpec == HTML 1.0?
 - Fear of the Tower of Babel



Rules for Contracts

- Don't make promises you can't keep...but don't hide power. [Lampson]
- There are no guarantees, ever.
 - Have a backup plan for what happens if “assurances” are not kept.
- Provide sufficient power to represent what promises the actor is explicitly NOT making.
 - E.g., temporary donation of resources
- Incorporate time: start/expiration time
 - Resource contracts are *leases* (or *tickets*).



Delegation

- Principle #3: Contracts enable delegation of powers.
 - Delegation is voluntary and provisional.
- It is a building block for creating useful concentrations of power in architecture.
 - Creates a potential for governance
 - Calendar scheduling, reservation
 - Double-edged sword?



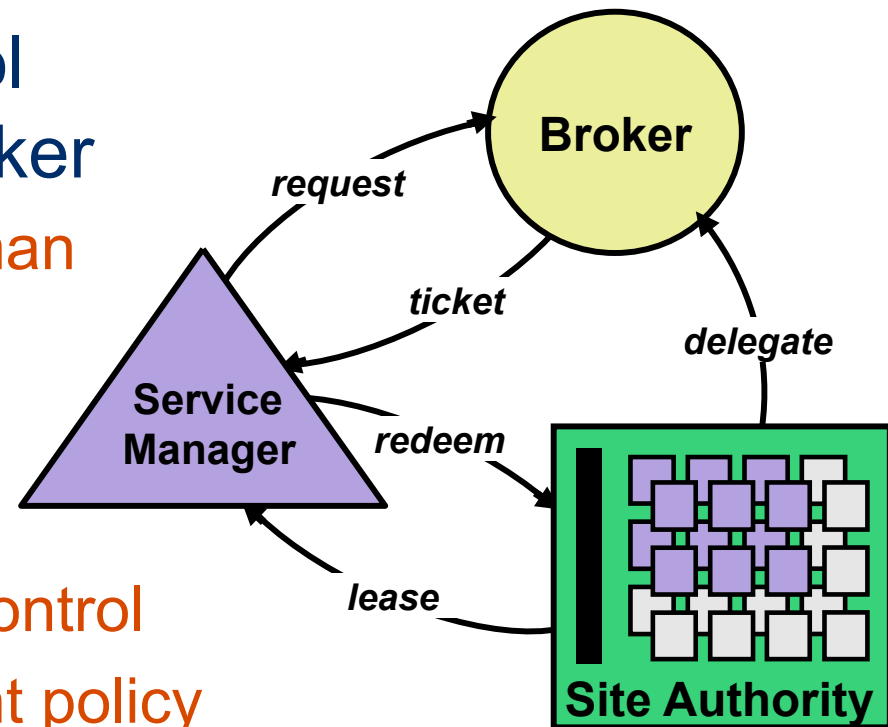
Aggregation

- Principle #4: aggregate the resources for a site or domain.
 - Primary interface is domain/site authority
- Abstraction/innovation boundary
 - Keep components simple
 - Placement/configuration flexibility for owner
 - Mask unscheduled outages by substitution
 - Leverage investment in technologies for site/domain management

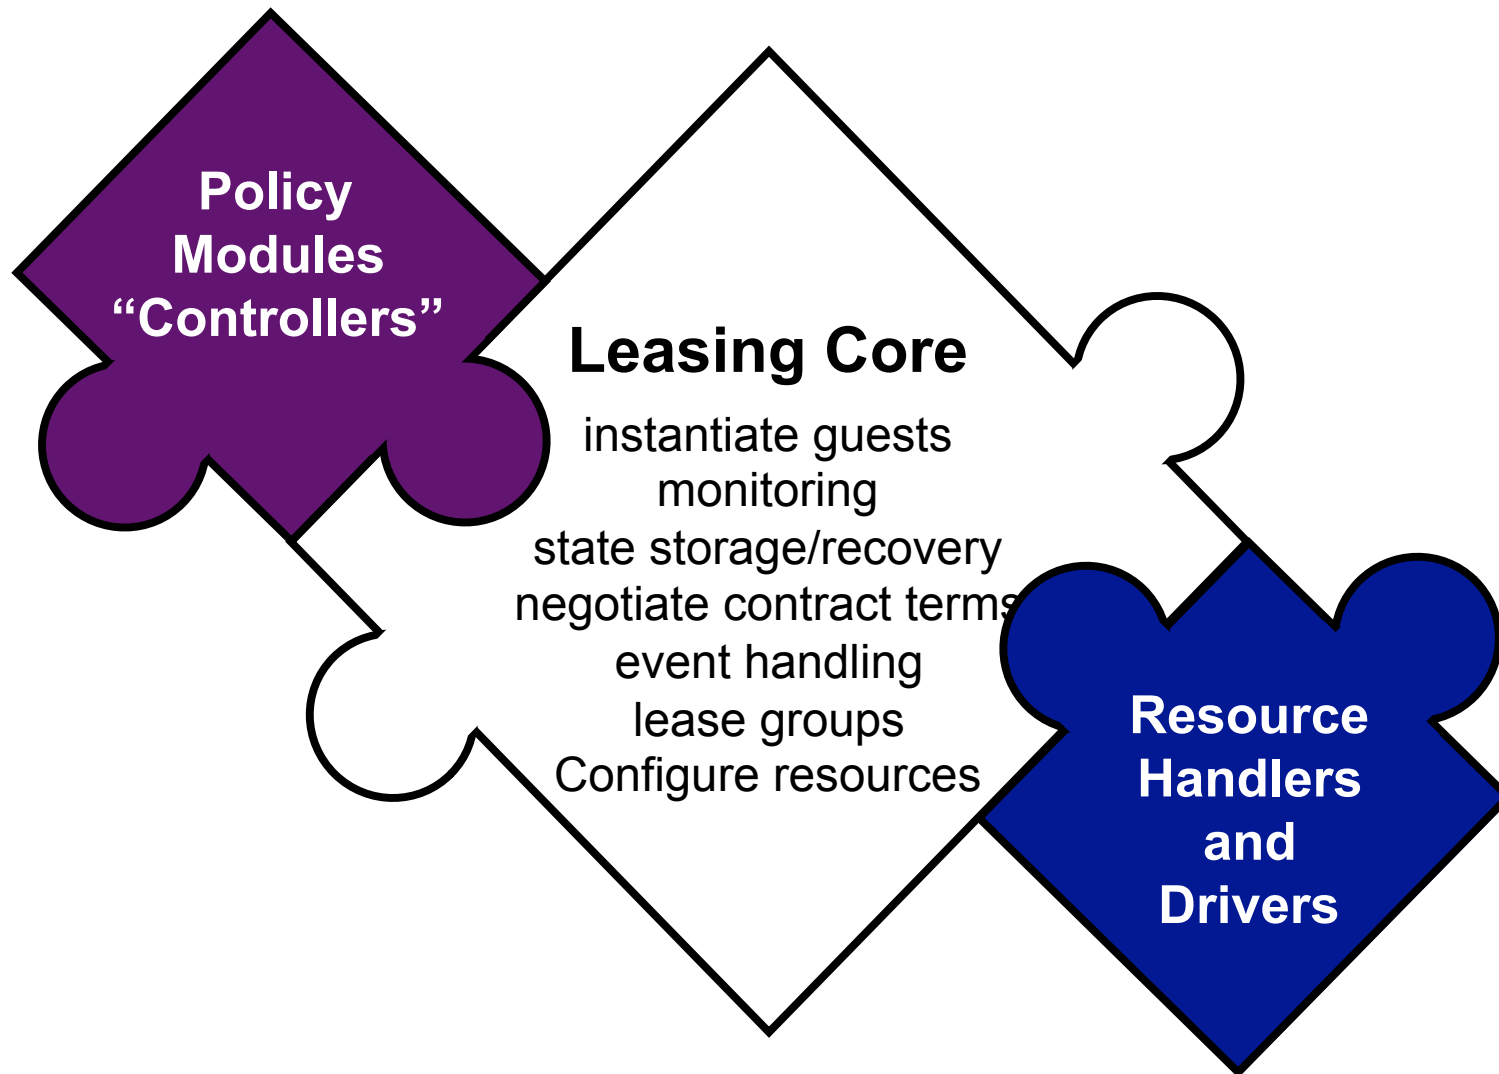


Brokers and Ticketing


- Sites delegate control of resources to a broker
 - Intermediary/middleman
- Factor allocation policy out of the site
 - Broker arbitrates resources under its control
 - Sites retain placement policy
- “Federation”
 - Site autonomy
 - Coordinated provisioning



Pluggable Resources and Policies



Orca: Actors and Protocols




GENI Principals



Researcher: A user that wishes to run an experiment or service in a slice, or a developer that provides a service used by other researchers.

Programmatic?
Automated?



A **slice authority (SA)** is responsible for the behavior of a set of slices, vouching for the users running experiments in each slice and taking appropriate action should the slice misbehave.

Who does it represent?
Policing or helper?



A **management authority (MA)** is responsible for some subset of substrate components: providing operational stability for those components, ensuring the components behave according to acceptable use policies, and executing the resource allocation wishes of the component owner.

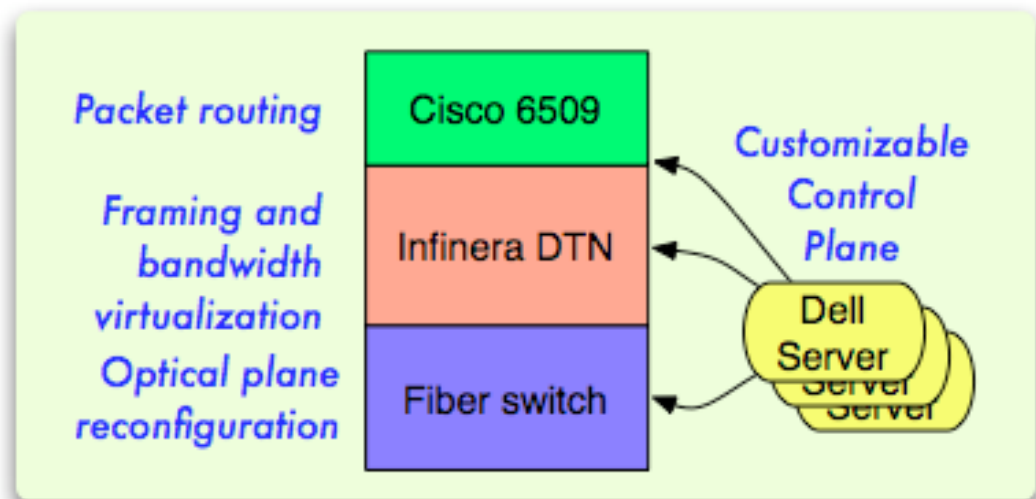
Who does it represent?
Policing or helper?

Aggregates? Clearinghouse? GSC Policies?



Breakable Experimental Network

- BEN is a metro network testbed jointly operated by researchers at Triangle Universities
- Designed to facilitate disruptive cross-layer research starting from the physical layer up
- 4 sites interconnected by dark fiber






Summary

- Factor actors/roles along the right boundaries.
 - stakeholders, innovation, tussle
- Open contracts with delegation
- Specific recommendations for GENI:
 - Aggregates are first-class entities
 - Component interface: permit innovation
 - Clearinghouse: enable policies under GSC direction



Thanks!





© Scott Adams, Inc./Dist. by UFS, Inc.

