Eric Boyd
Senior Director, Strategic Projects, Internet2
Ed Balas
Manager, Software Engineering, Indiana University

ENABLING INNOVATION THROUGH NETWORK VIRTUALIZATION
(AND INTEGRATION OF COMPUTE AND STORAGE)


We the People
INTERNET2®
Unleashing new waves of global discovery, together.

Advanced Layer2 Service Deployment

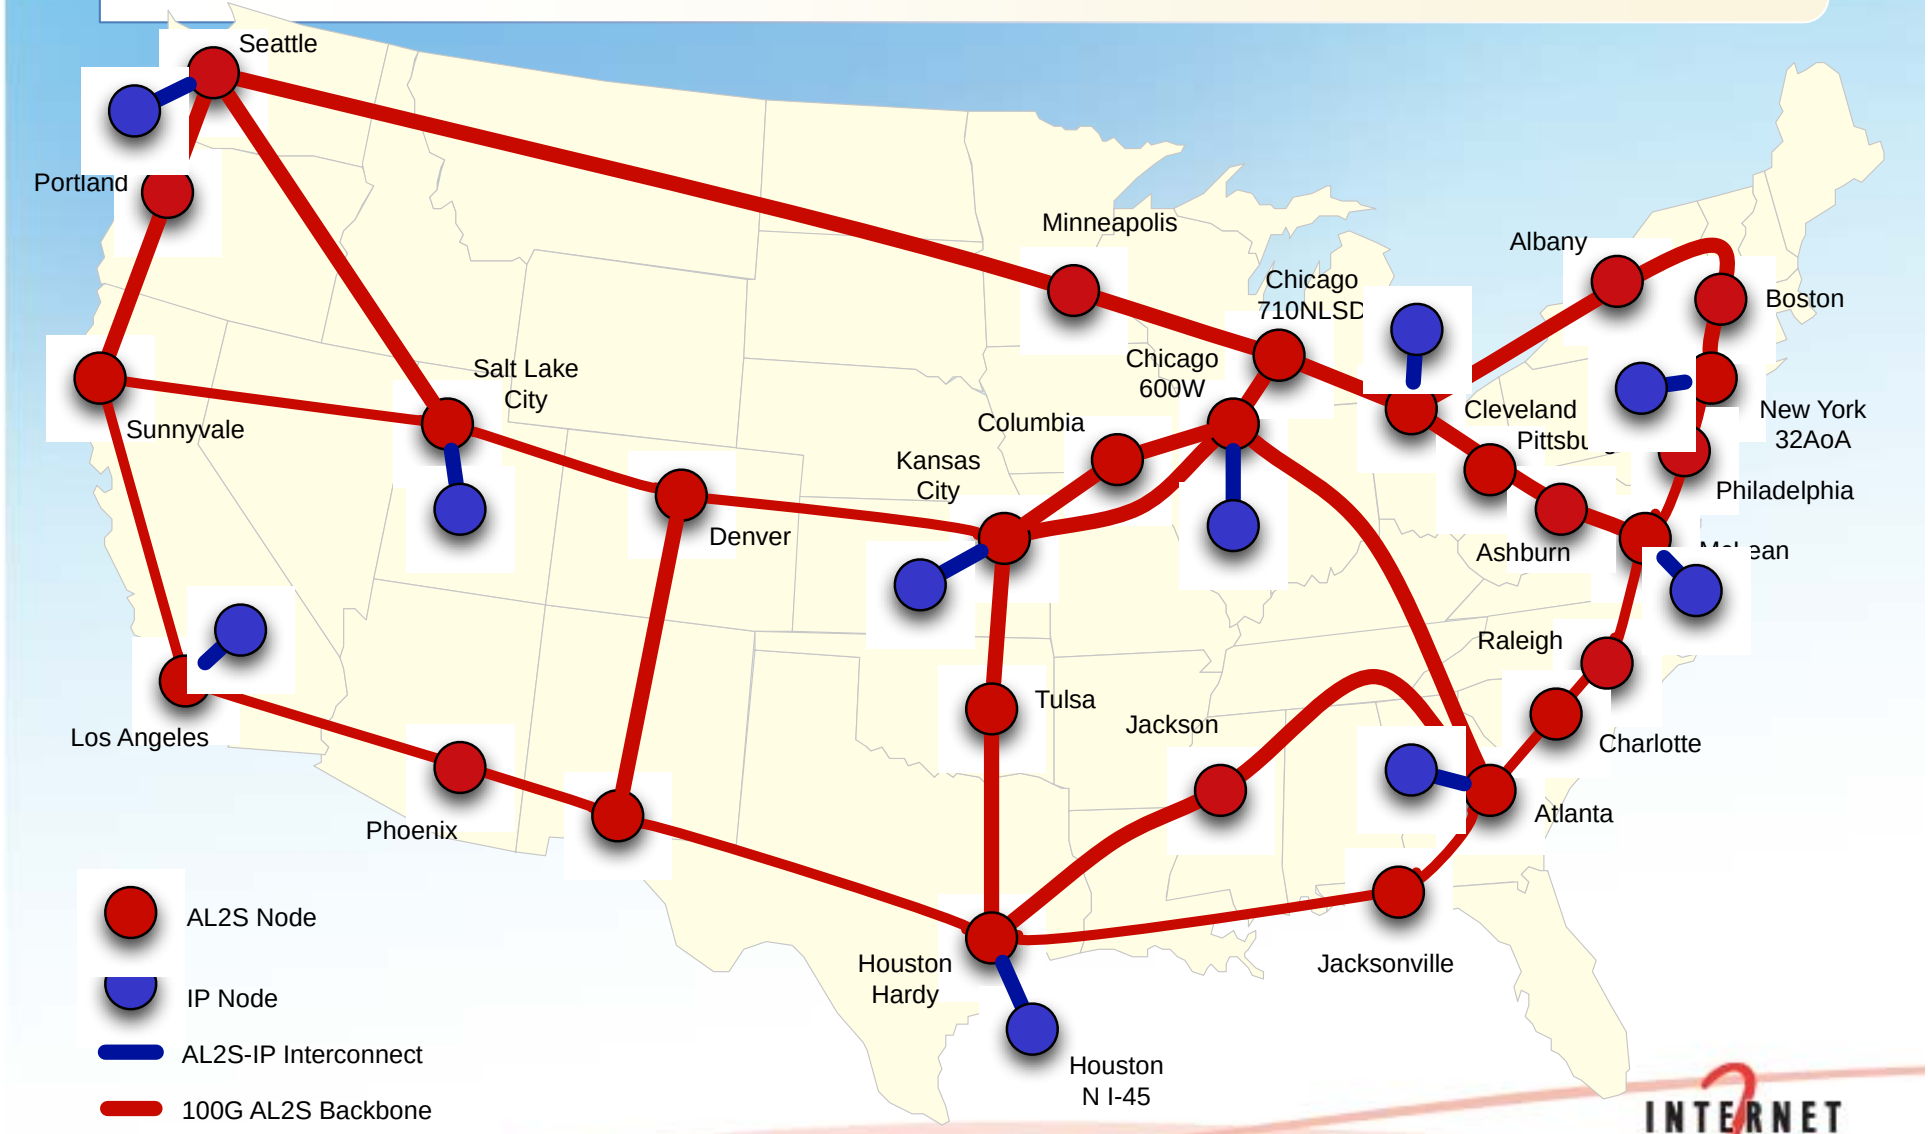

This is what we have been able to say for about a year:
The **100G** testbed of innovation for tomorrow's Internet is available nationwide, right now.

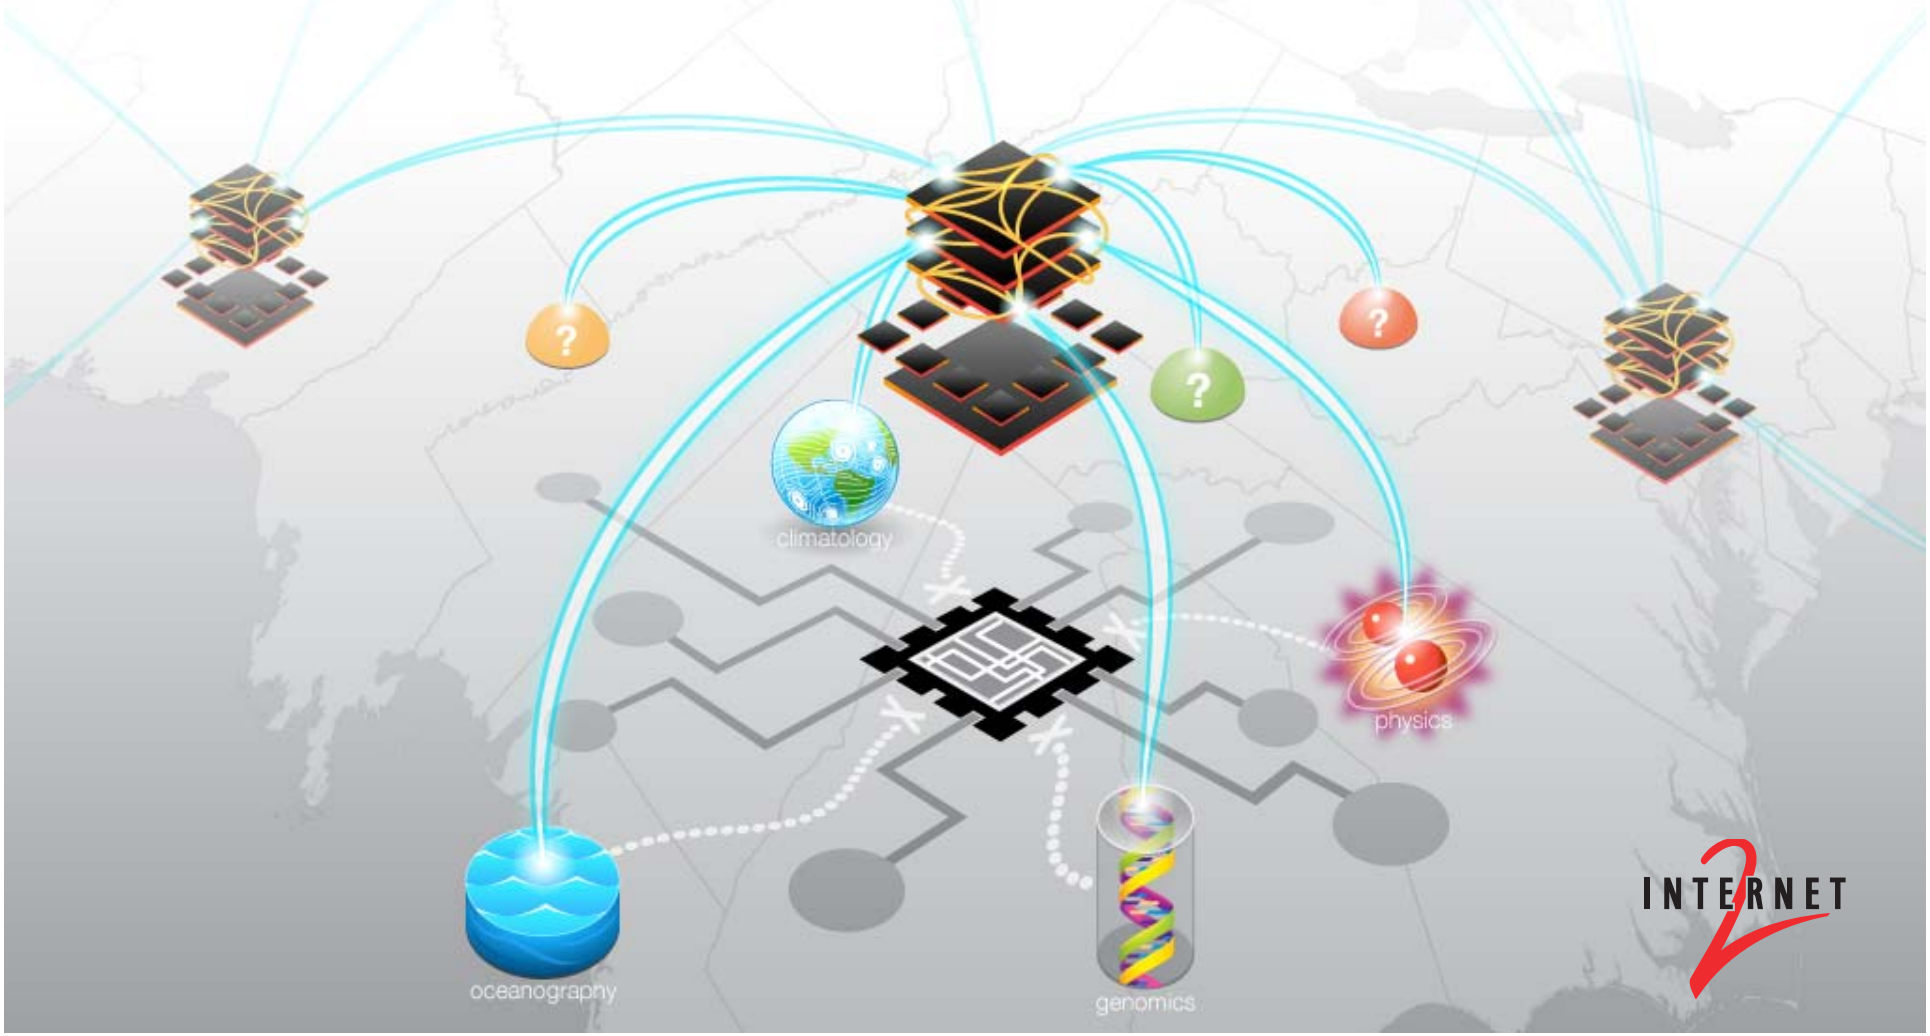

Does this create a platform for innovation?

Abundant bandwidth to enable innovation?

Programmability to encourage application innovation?

Support data intensive science?

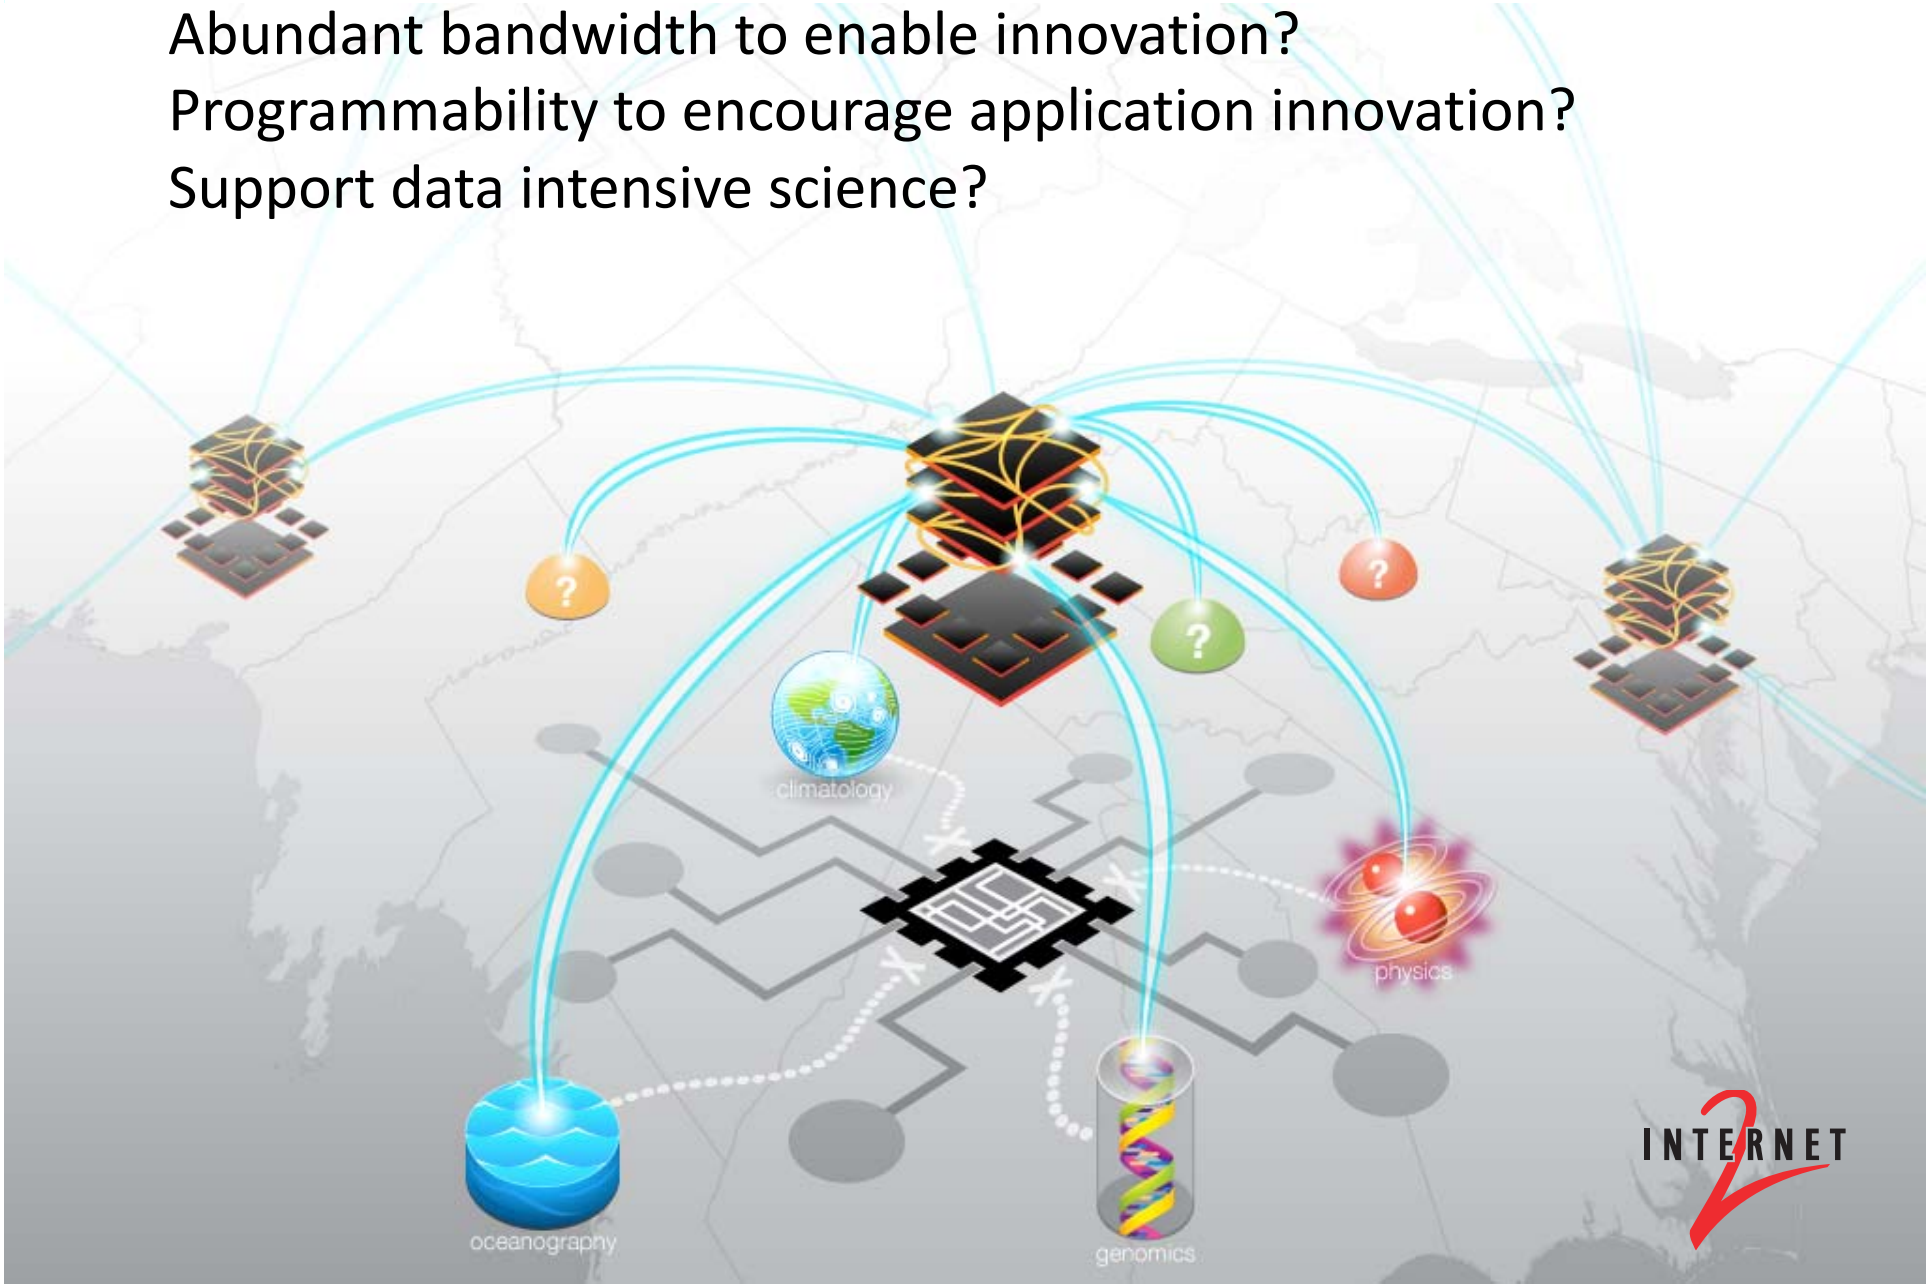

Does this create a platform for innovation?

Abundant bandwidth to enable innovation? ☑

Programmability to encourage application innovation? ☑

Support data intensive science? ☑

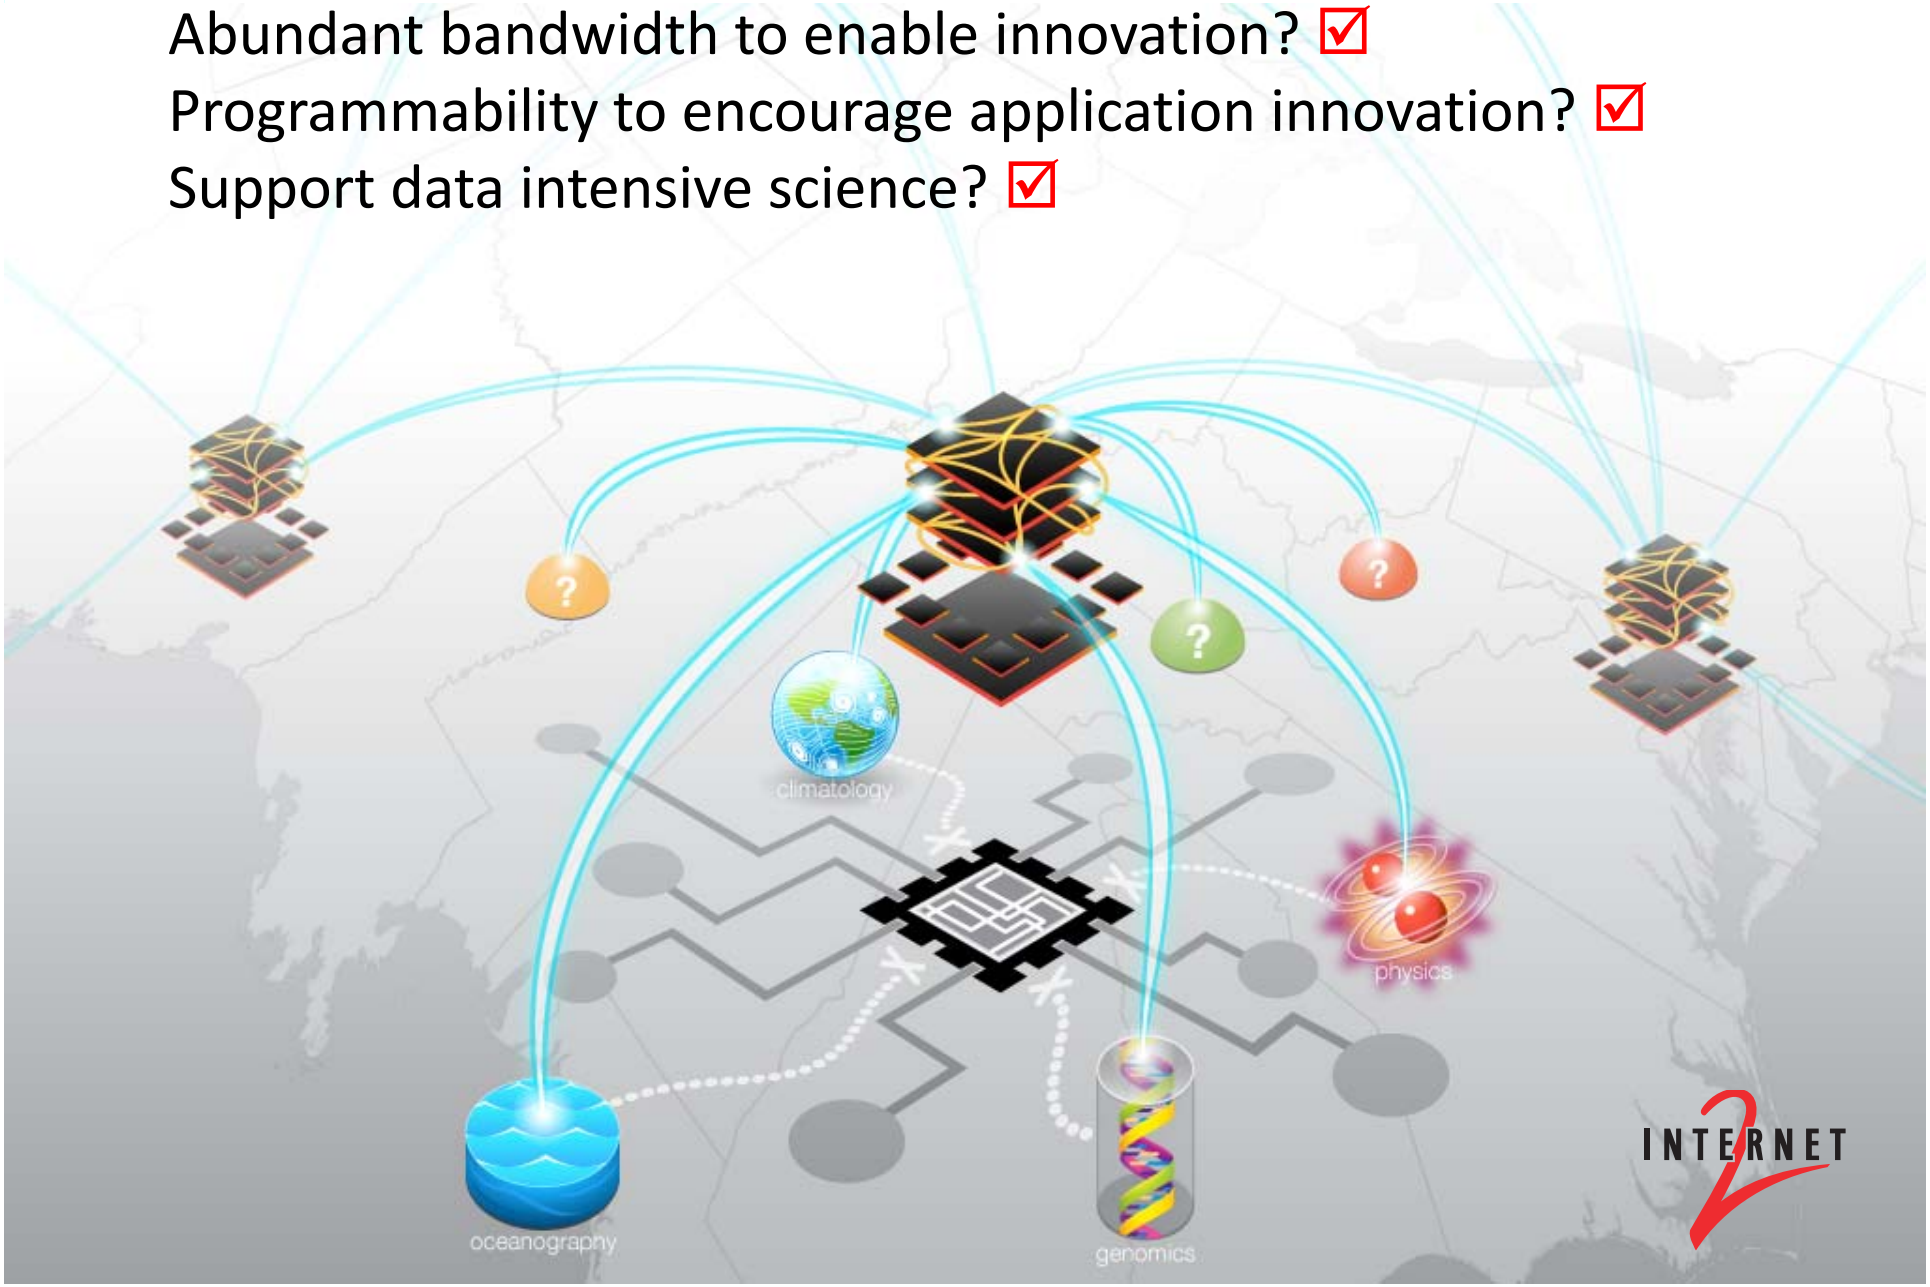

TODAY

Does this create a platform for innovation?

Abundant bandwidth to enable innovation? ☑

Programmability to encourage application innovation? ☑

Support data intensive science? ☑

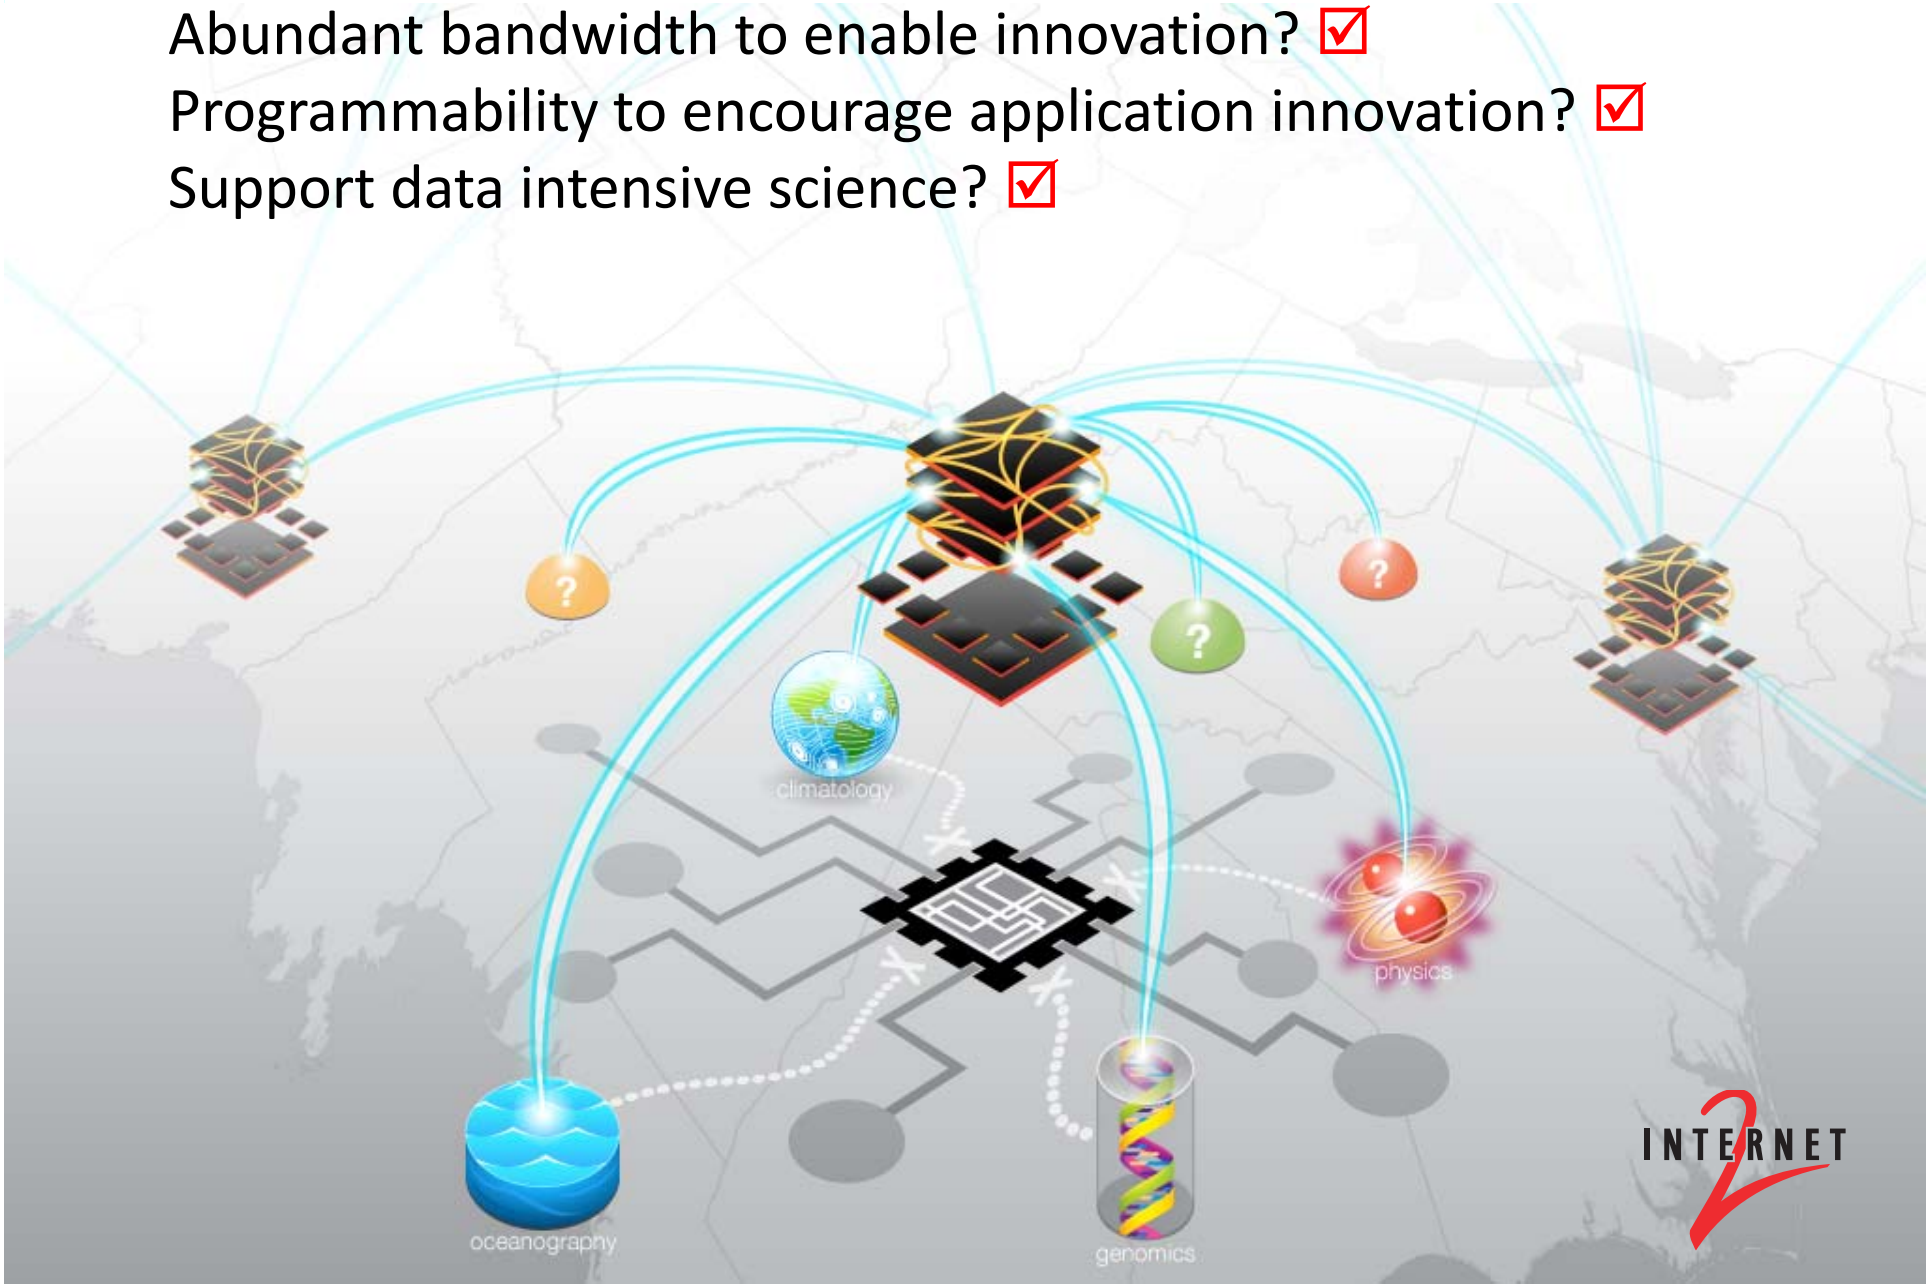

Does this create a platform for innovation?

- Abundant bandwidth to enable innovation? ☑
- Software-defined networking substrate? ☑
- Support data intensive science? ☑
- Virtualization? ☐
- Integrate network with compute and storage? ☐

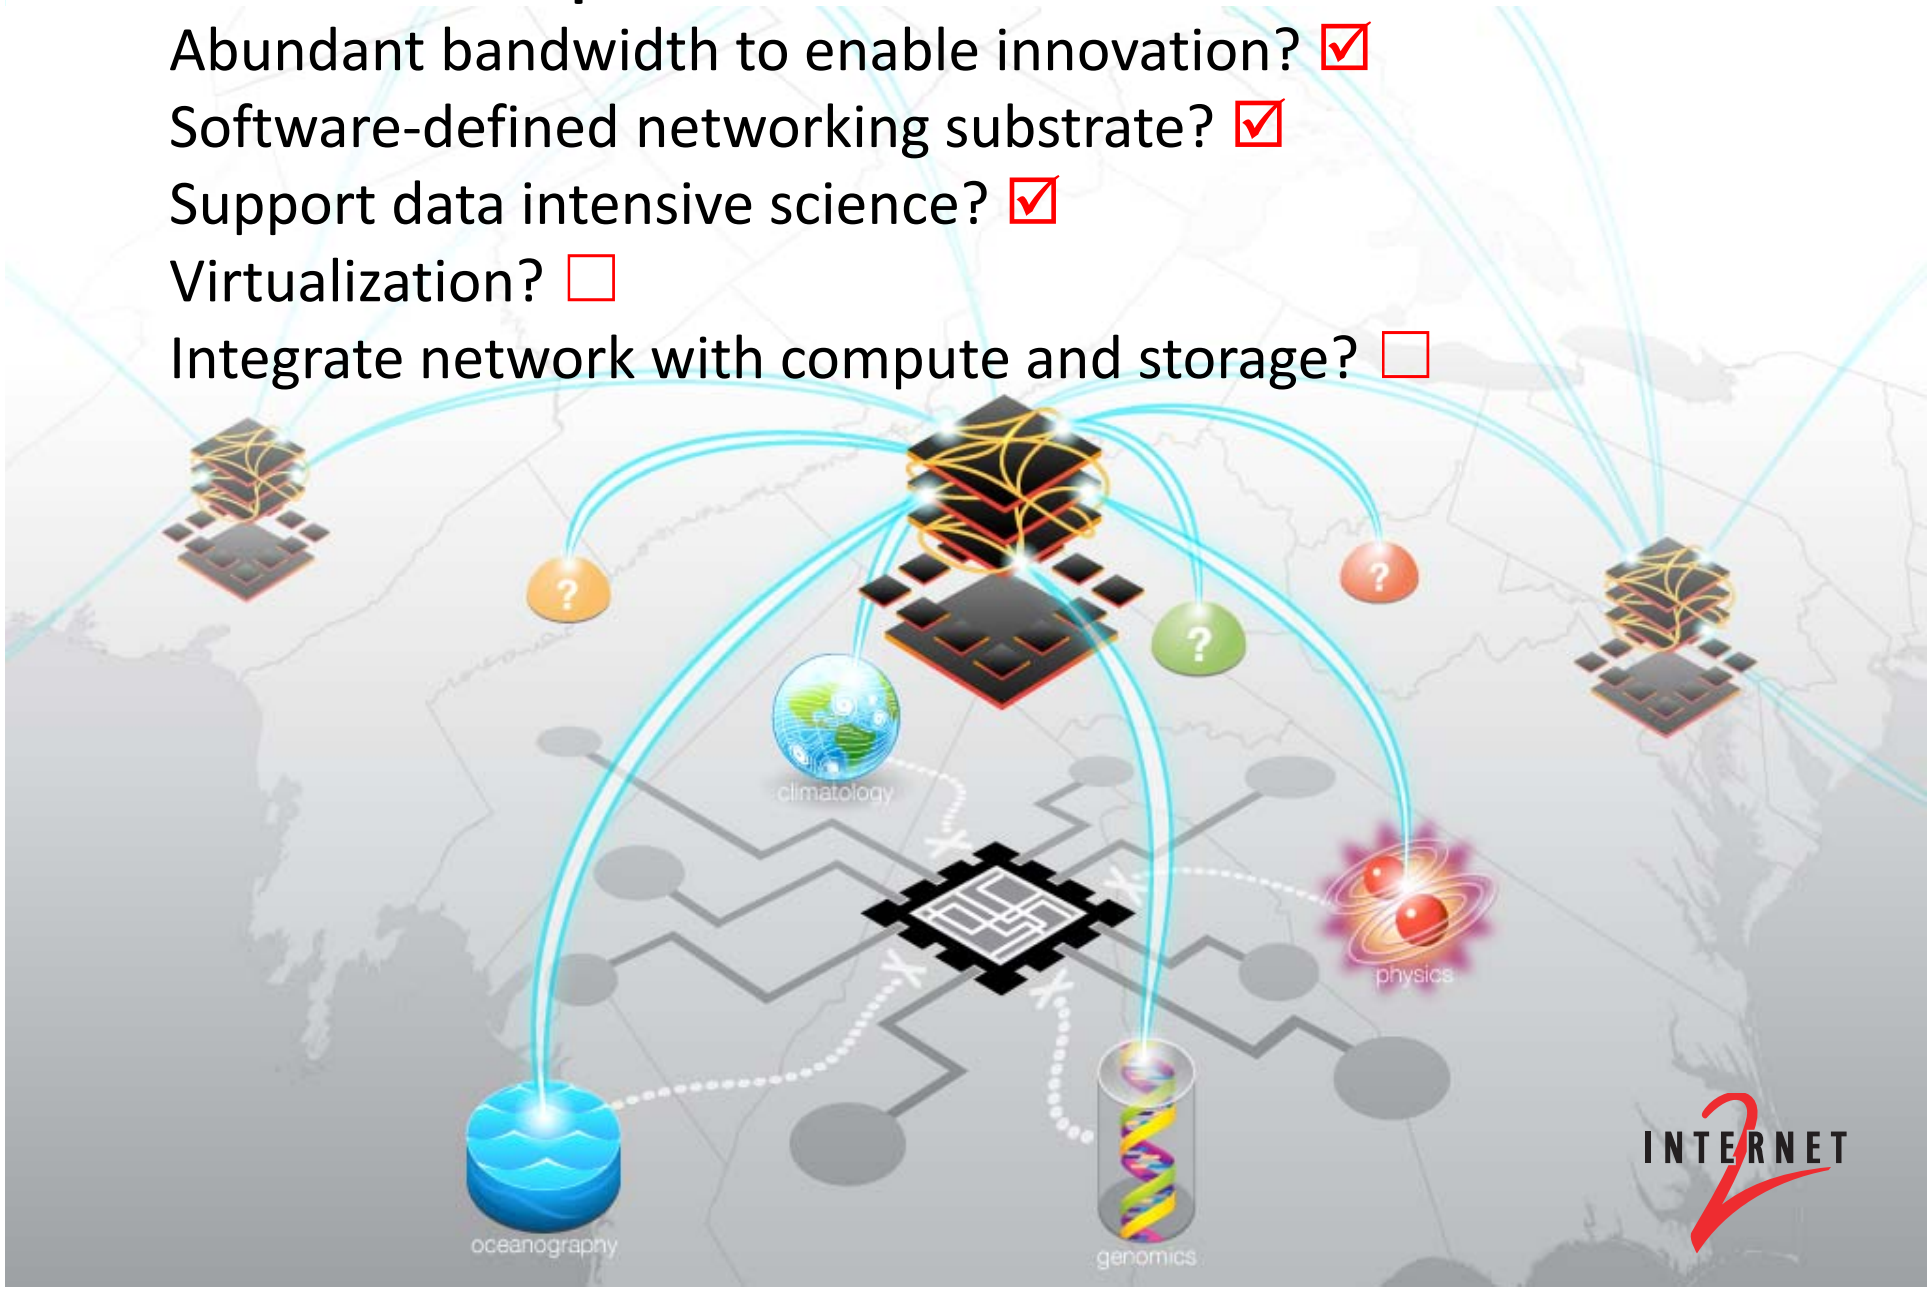

TOMORROW

Does this create a platform for innovation?

- Abundant bandwidth to enable innovation? ☑
- Software-defined networking substrate? ☑
- Support data intensive science? ☑
- *Virtualization*? ☐
- Integrate network with compute and storage? ☐

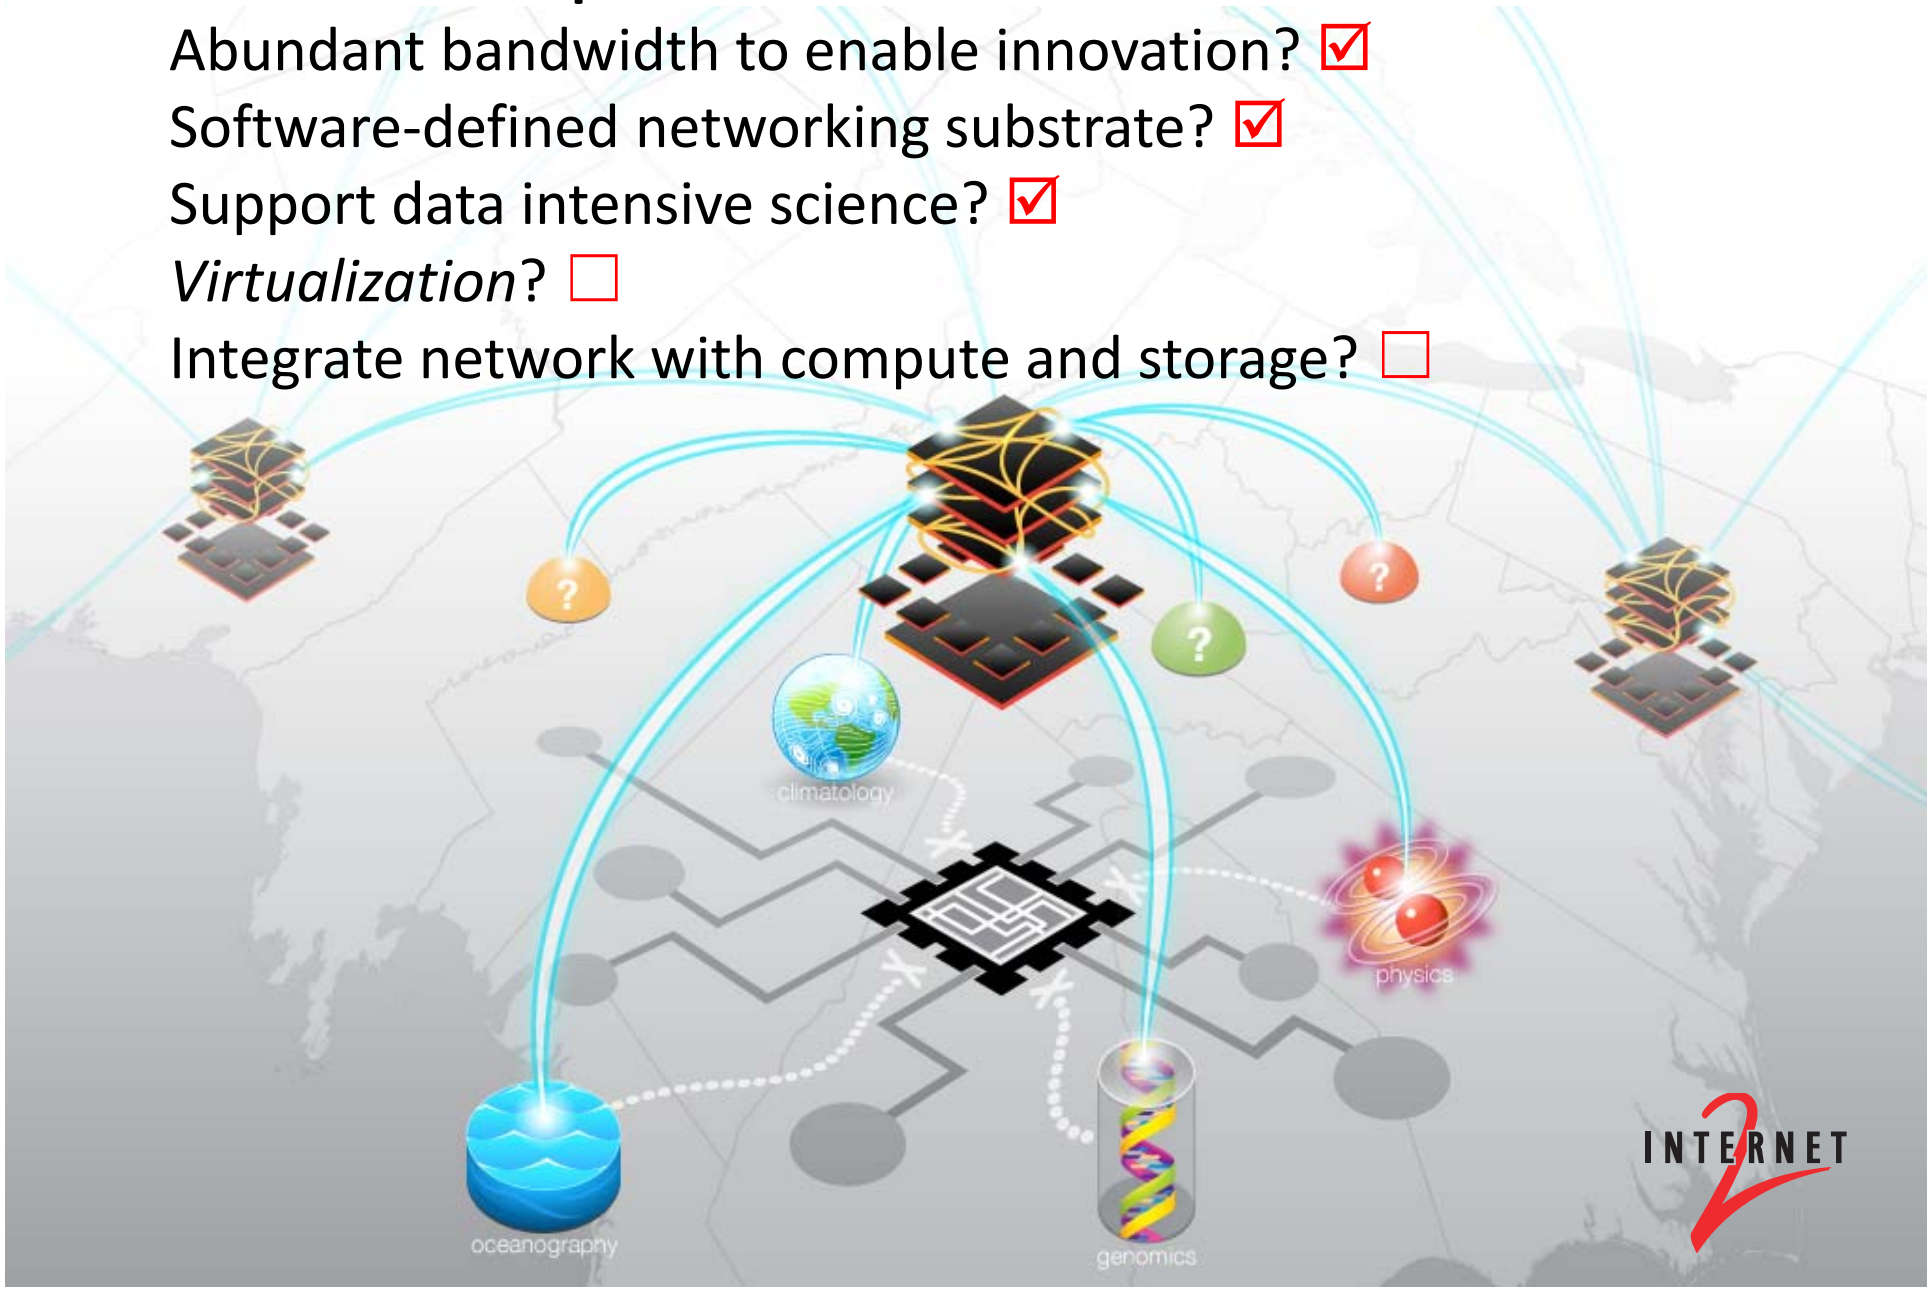

So what does this mean for GLIF Tech?

- We have a great deal of innovation, experimentation, and deployment in areas such as:
  - 100G networking around the globe
  - Software-defined networking
  - Enabling big science flows
- We need understand the various approaches to:
  - Network Virtualization
  - Integration with compute and storage
- Questions to ponder:
  - Does the GLIF community need a common approach to network virtualization?
  - Does the GLIF community need an integrated approach to network virtualization?
  - **More importantly: How does the networking community integrate itself into the compute and storage community?**

INTERNET 2

2013 Internet2 Innovative Application Awards

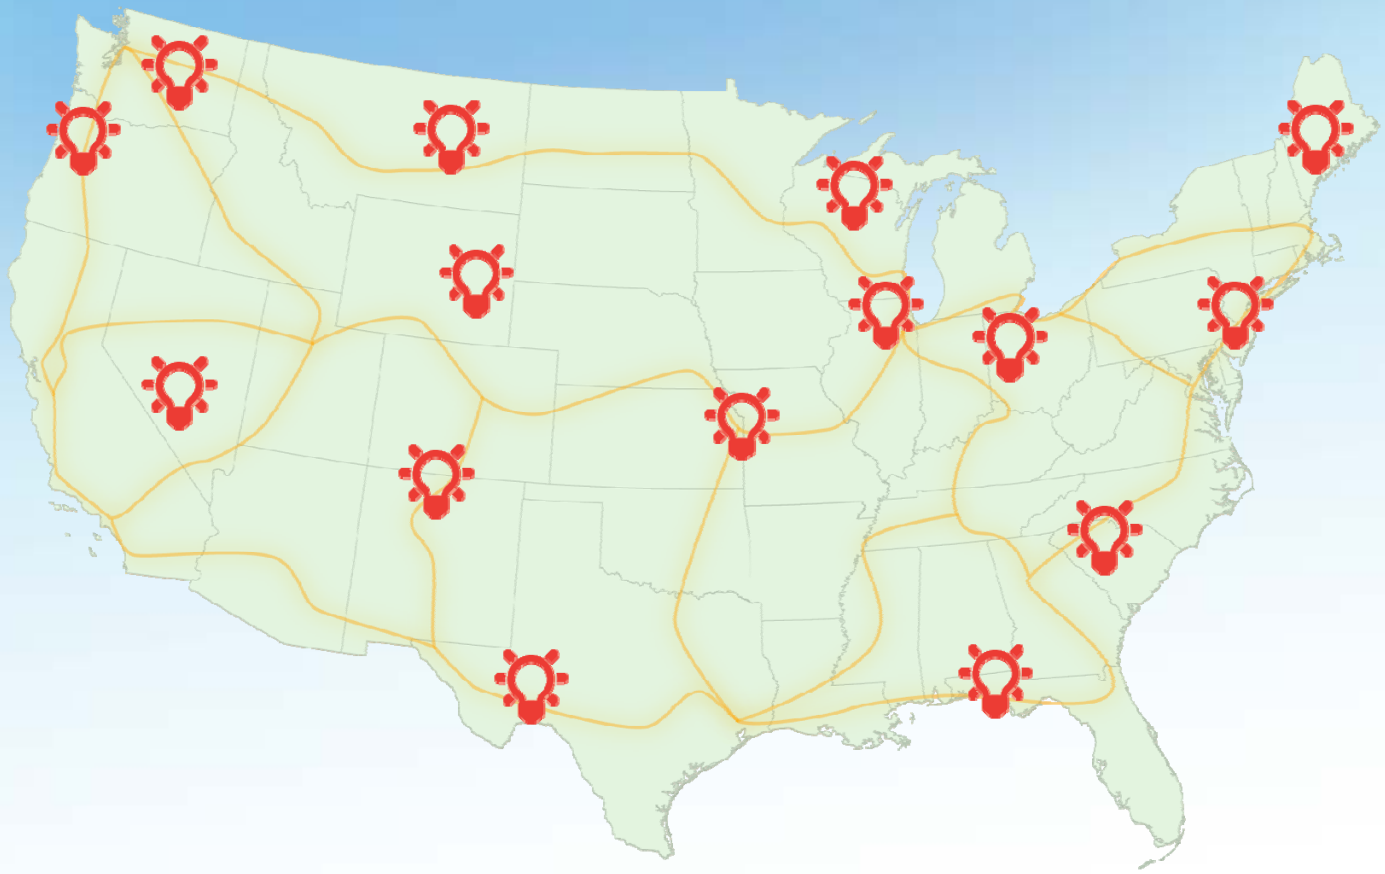

GOLD

Juniper Networks

SILVER

ciena

BRONZE

Brocade

Internet2

2013 Internet2 Innovative Application Awards

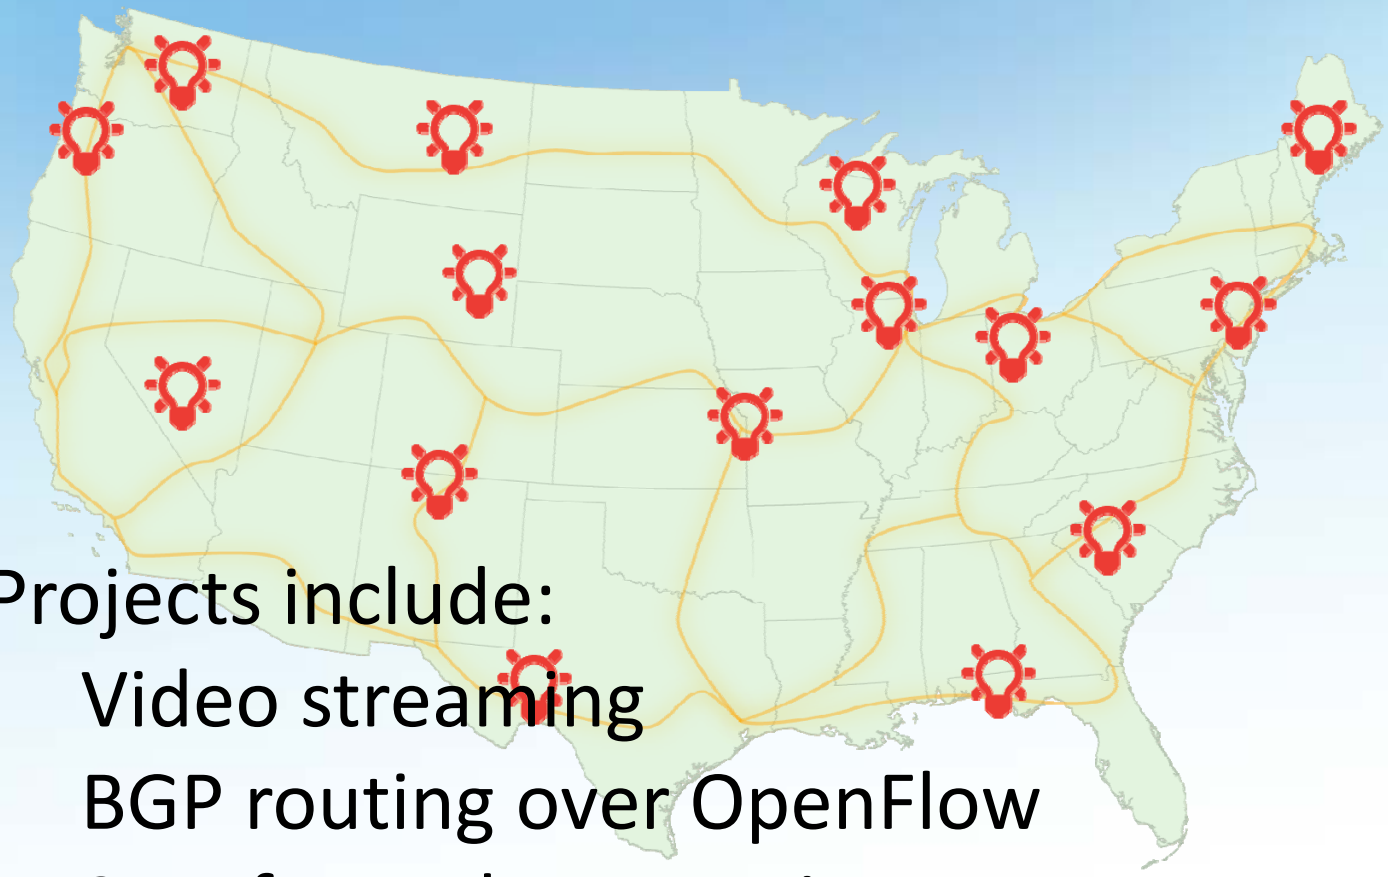

Projects include:

- Video streaming
- BGP routing over OpenFlow
- SDN for Exchange Points
- GENI integration

GOLD

Juniper Networks

SILVER

ciena

BRONZE

Brocade

Internet2

- We've had virtualization of storage and servers for quite some time
- How to define Network Virtualization?
- "Virtualization is the core principle in overlays, both allowing nodes to treat an overlay as if it were the native network, and allowing multiple overlays to simultaneously use the same underlying overlay infrastructure." (2004 – Anderson, Peterson, Shenker, Turner)

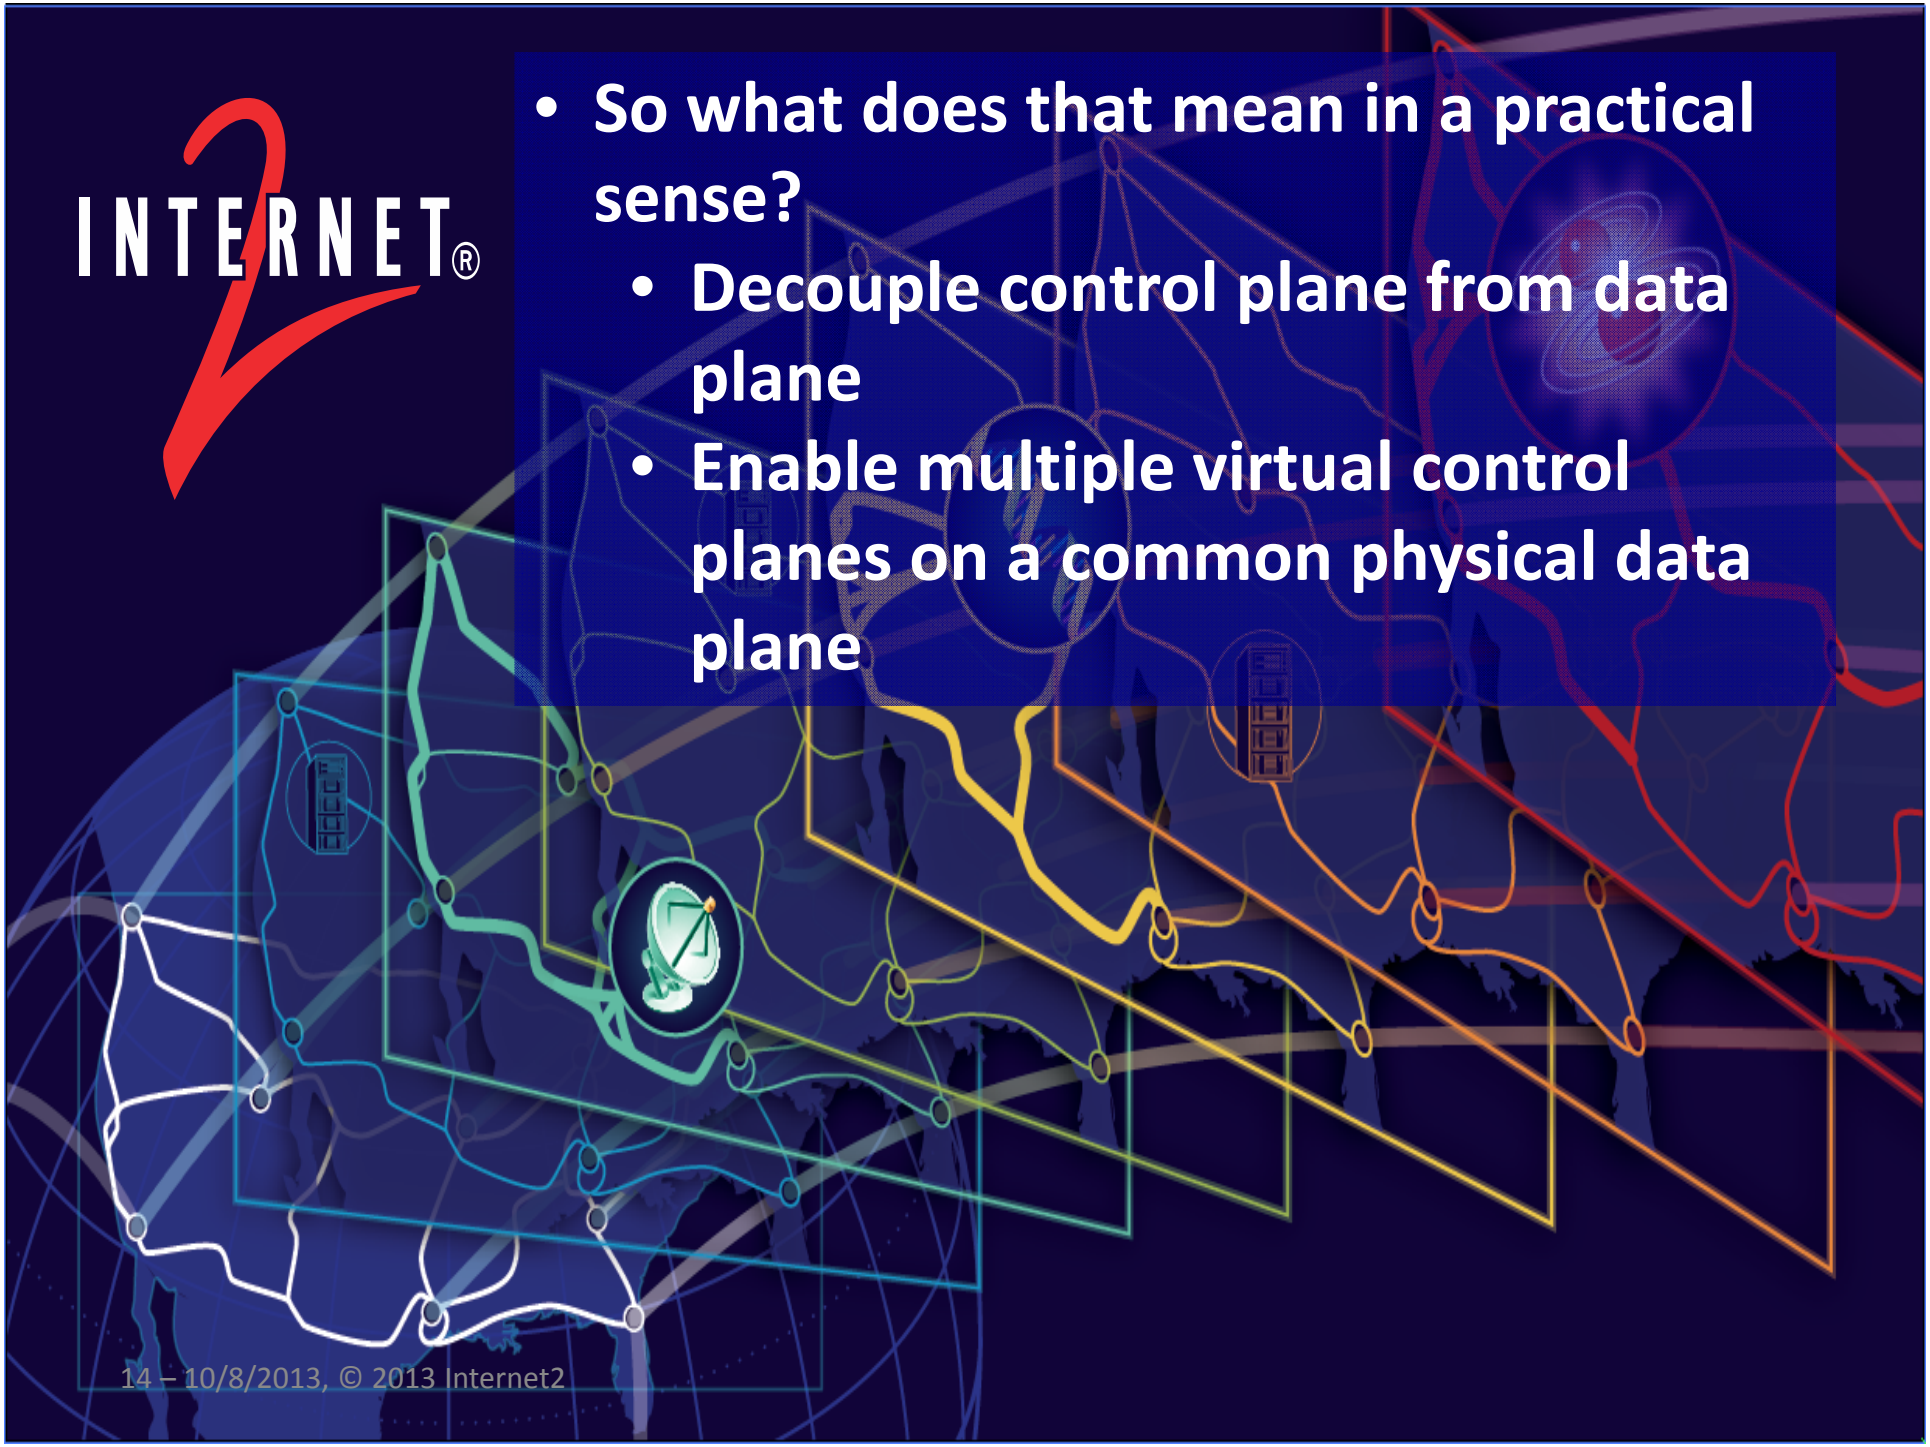

- So what does that mean in a practical sense?
  - Decouple control plane from data plane
  - Enable multiple virtual control planes on a common physical data plane

- Internet2 in partnership with Indiana University has been building / deploying an architecture to support network virtualization
  - Provide network multi-tenancy at Layer 2 and Layer 3
  - Enforce non-overlapping Layer 2 tag-based flowspace
    - Experiment Foo can use VLAN tag range 1-200 (a sliver)
    - Experiment Bar can use VLAN tag range 201-400 (a sliver)

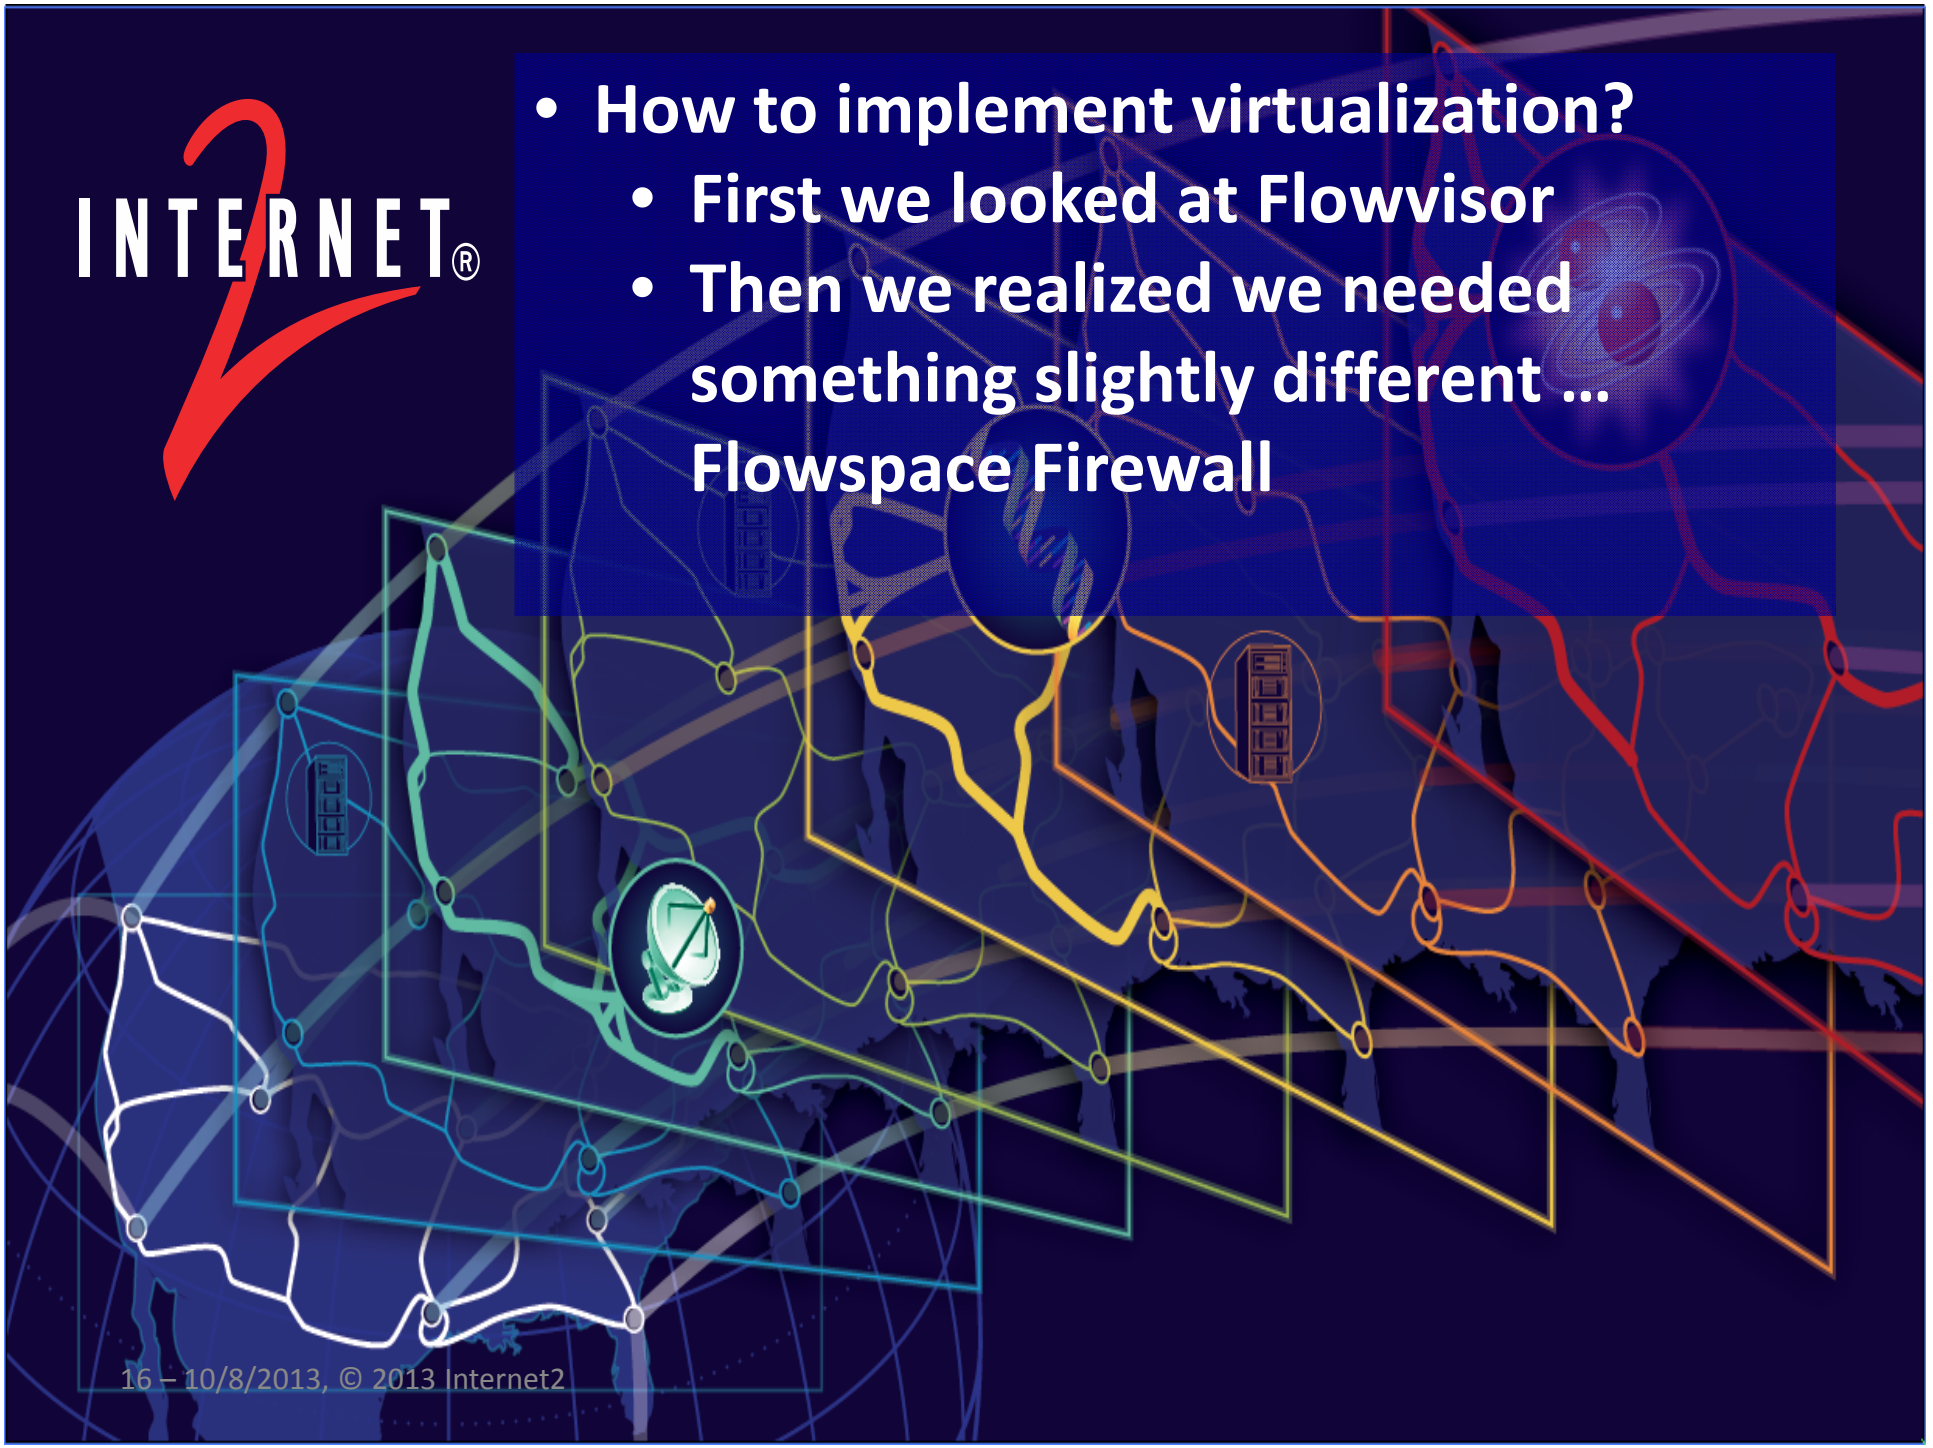
INTERNET®
• How to implement virtualization?
• First we looked at Flowvisor
• Then we realized we needed something slightly different … Flowspace Firewall

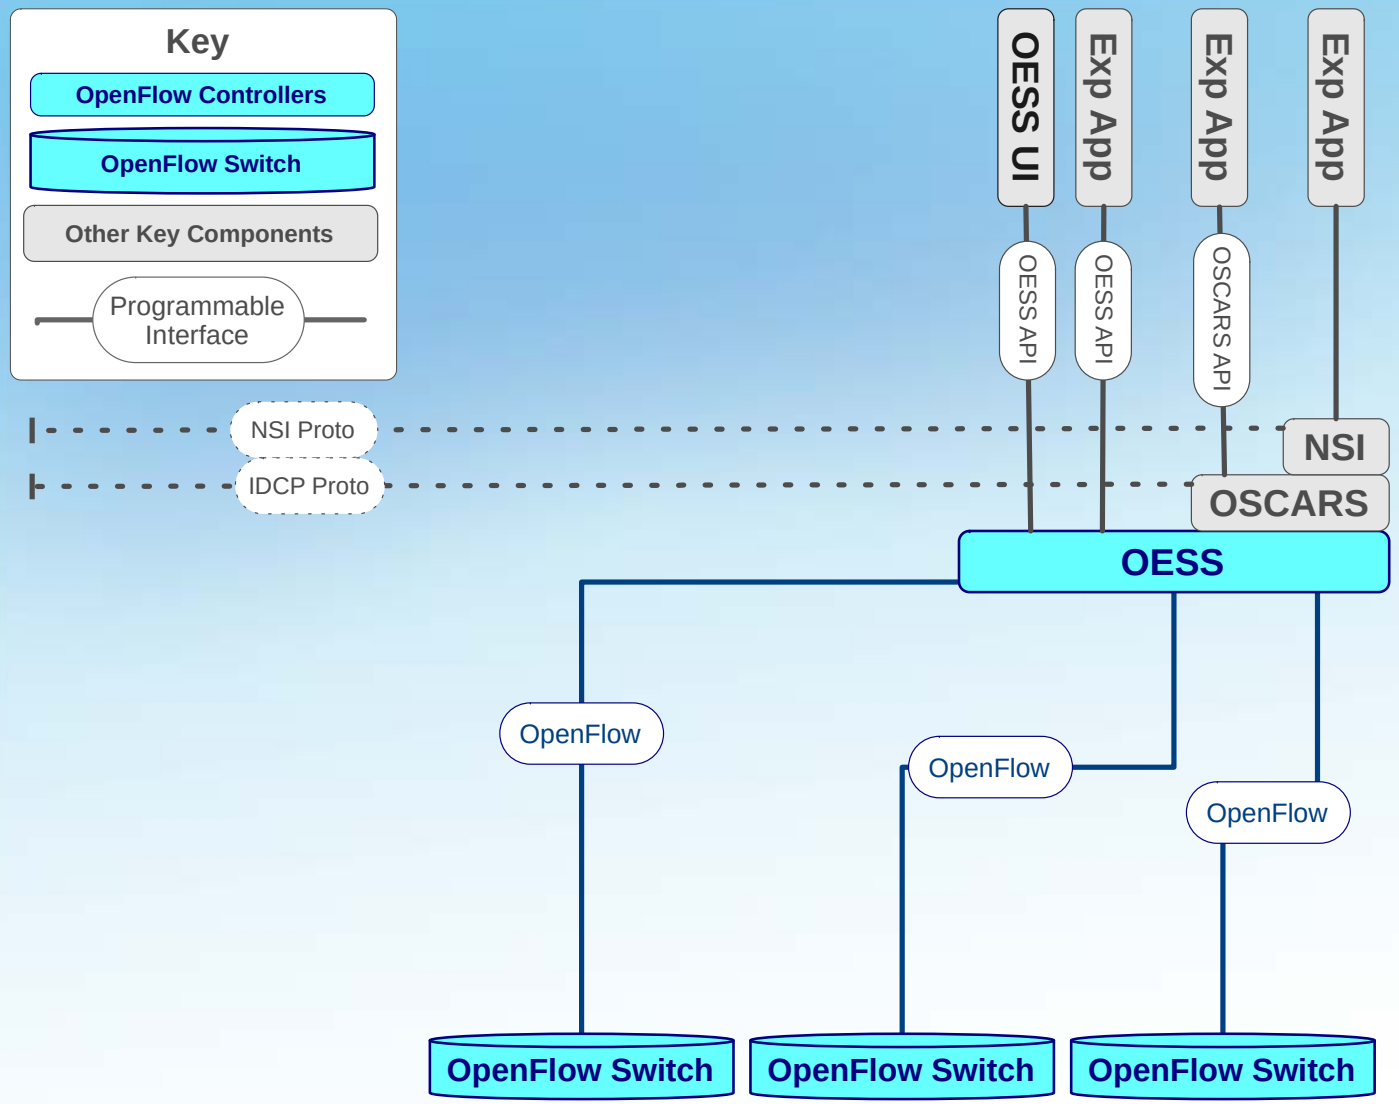

Current Configuration

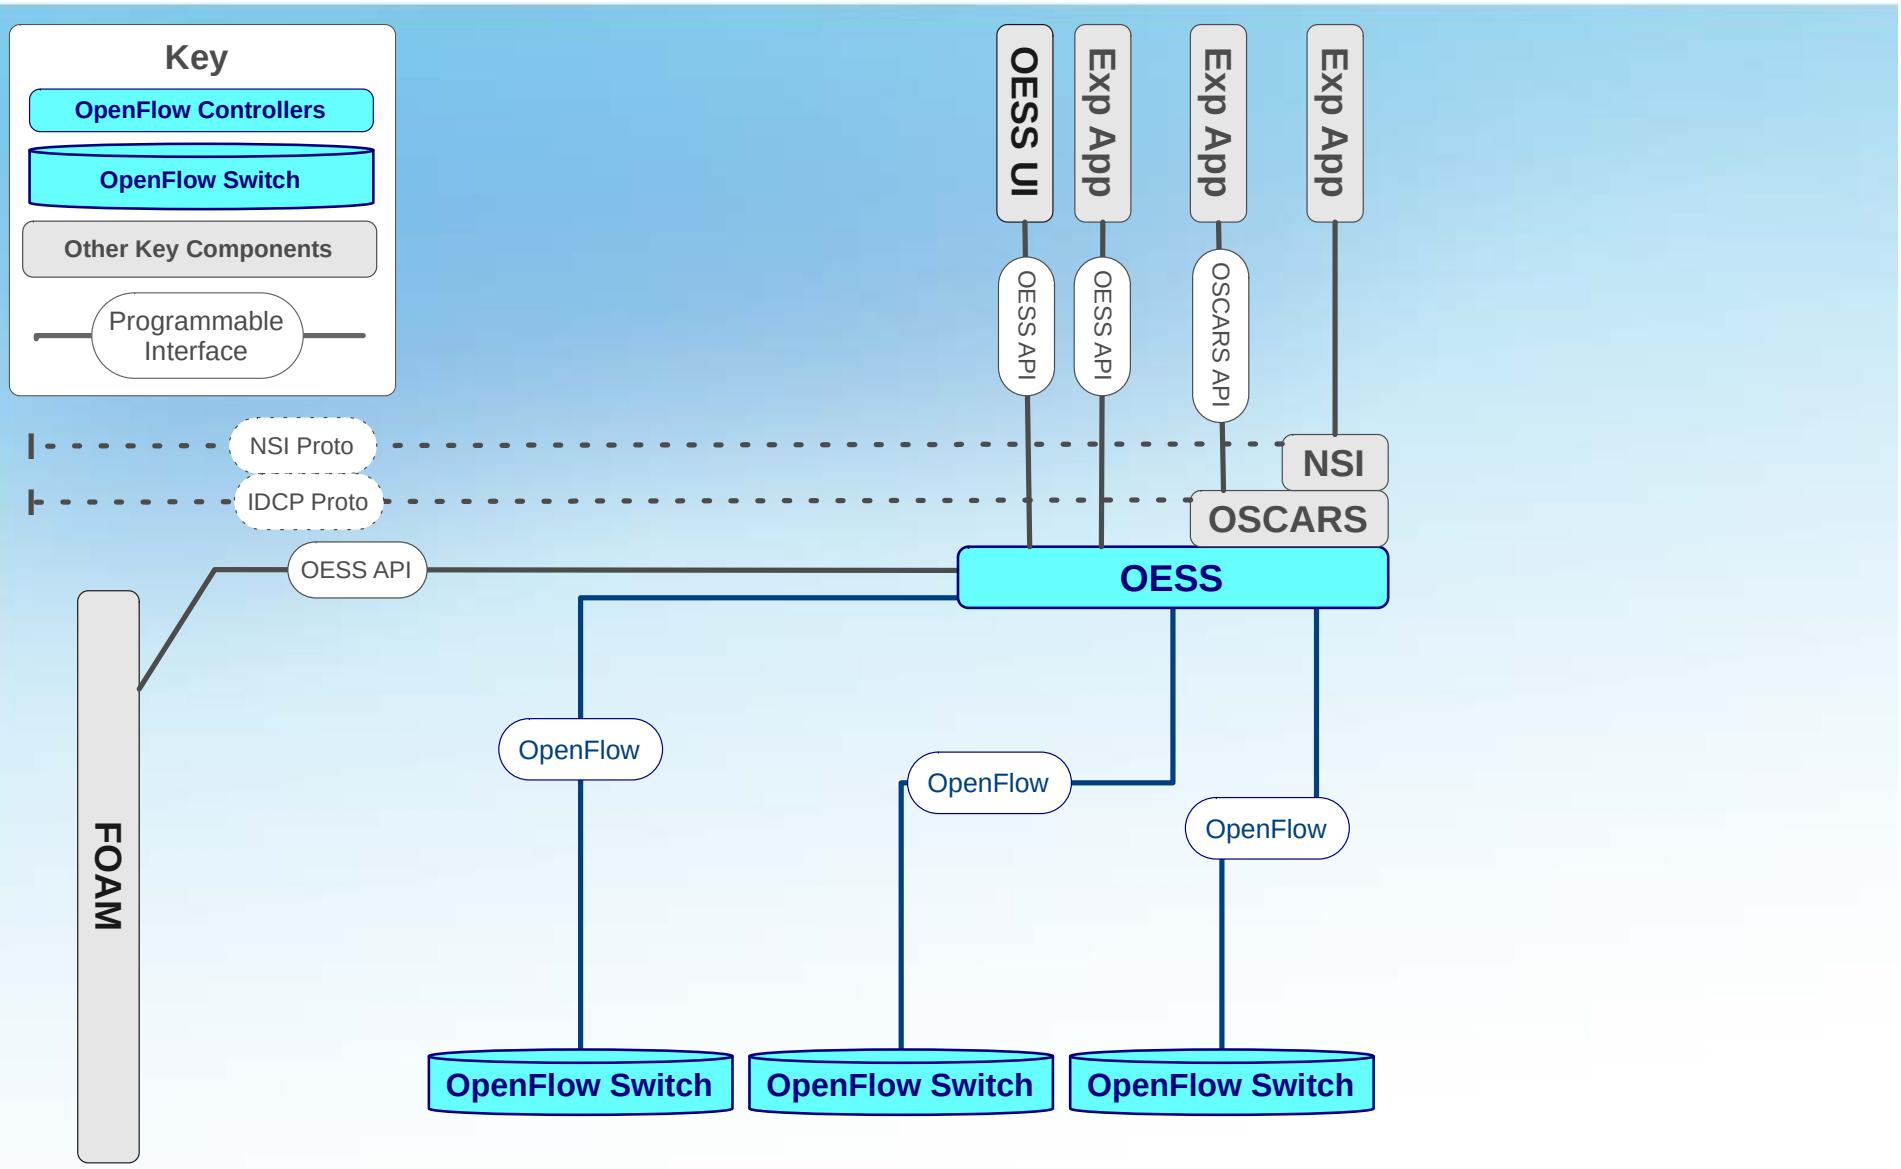

AL2S Software Stack Early Q4 2013

INTERNET 2

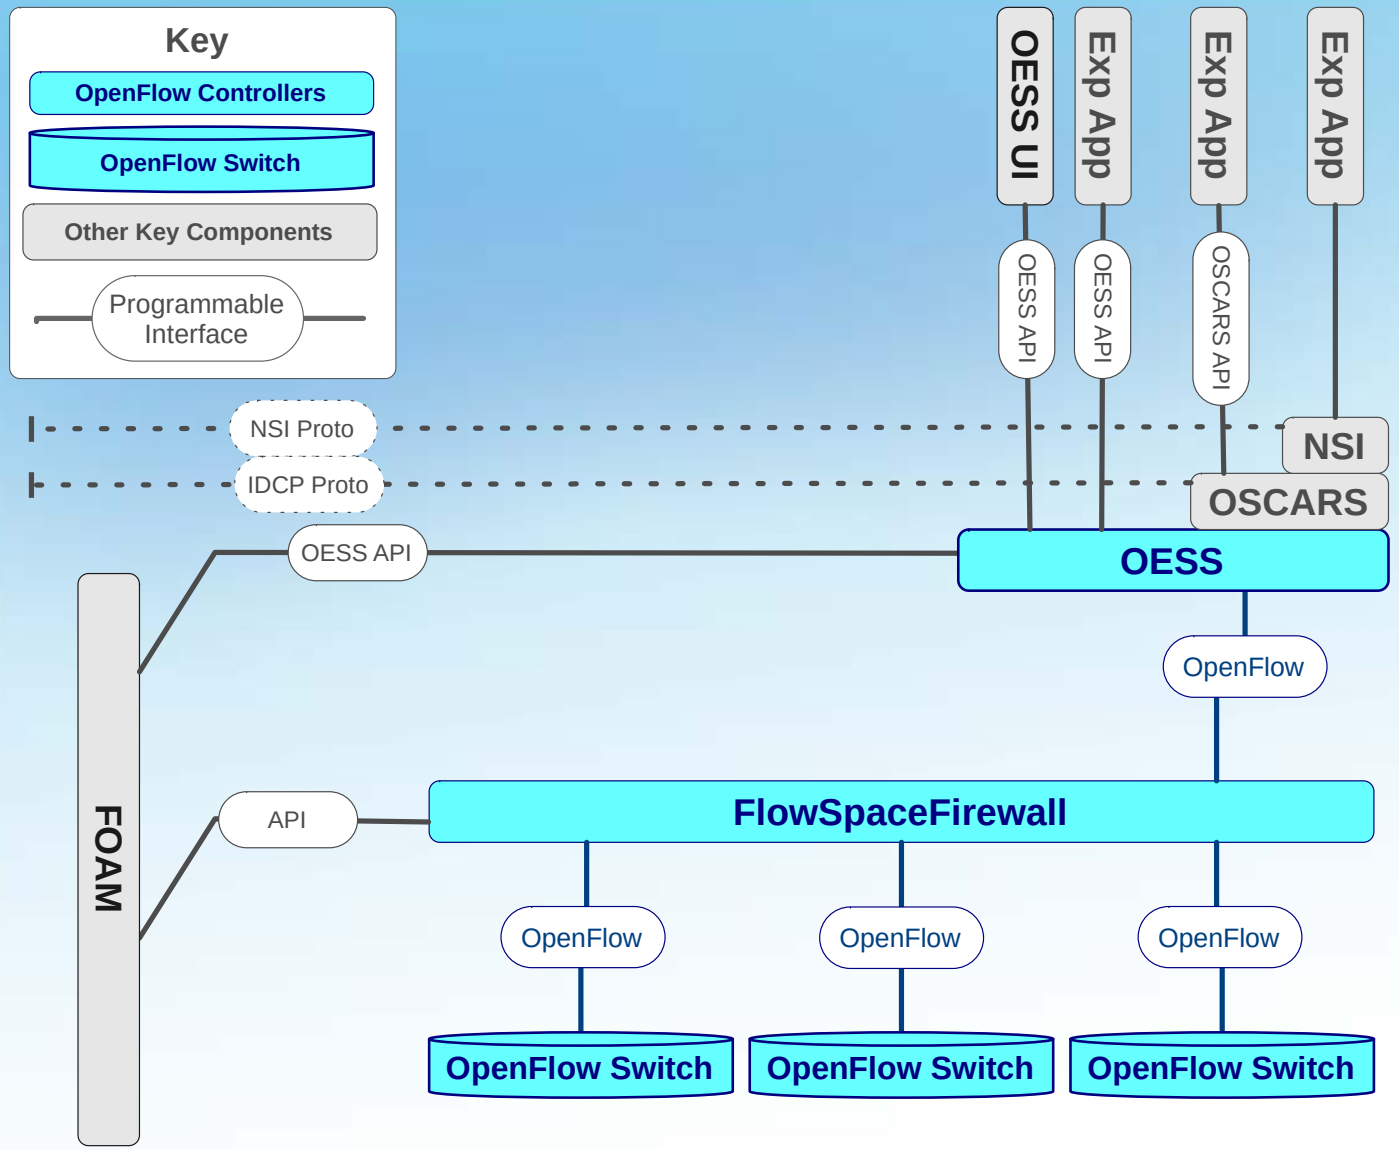

AL2S Software Stack Late Q4 2013

INTERNET 2

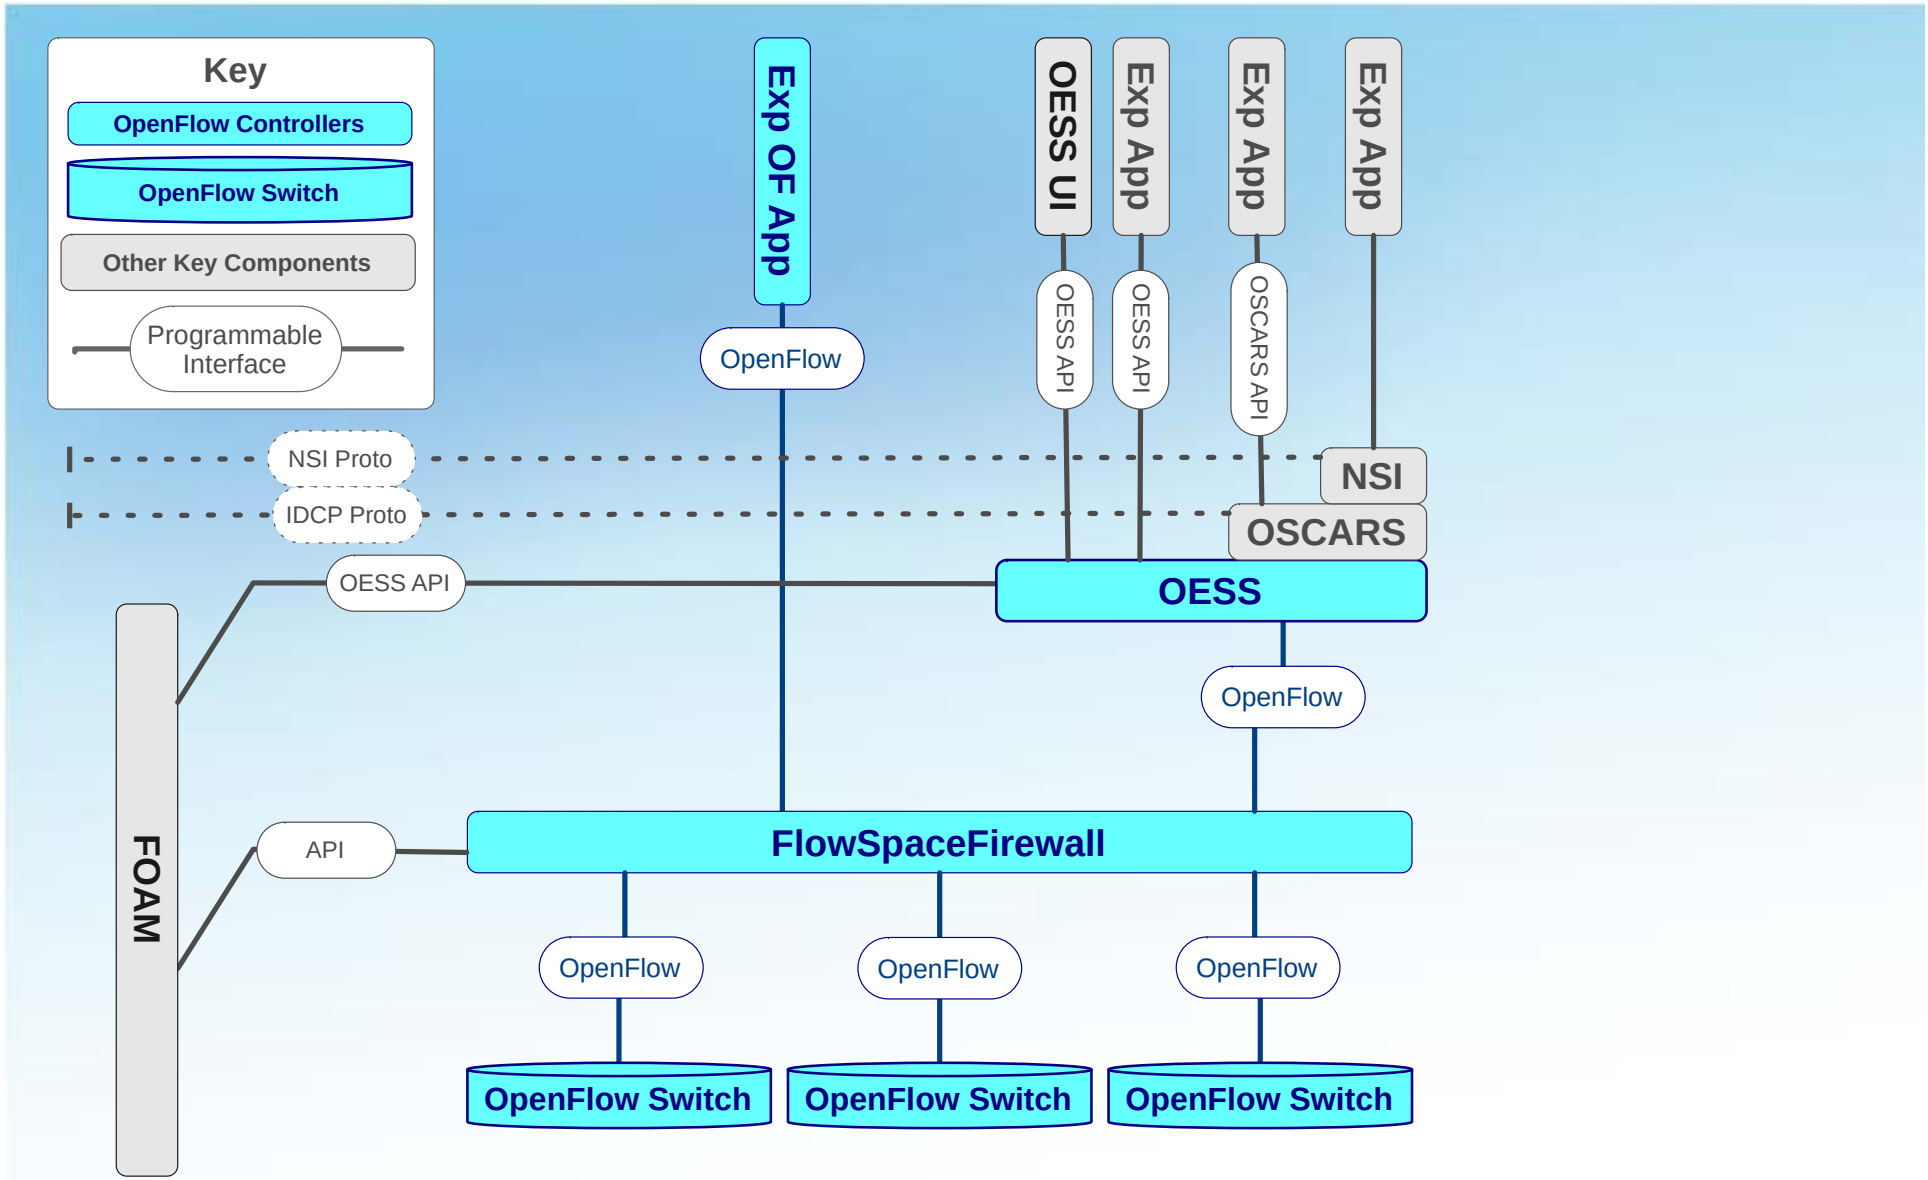

AL2S Software Stack Q1 Early 2014

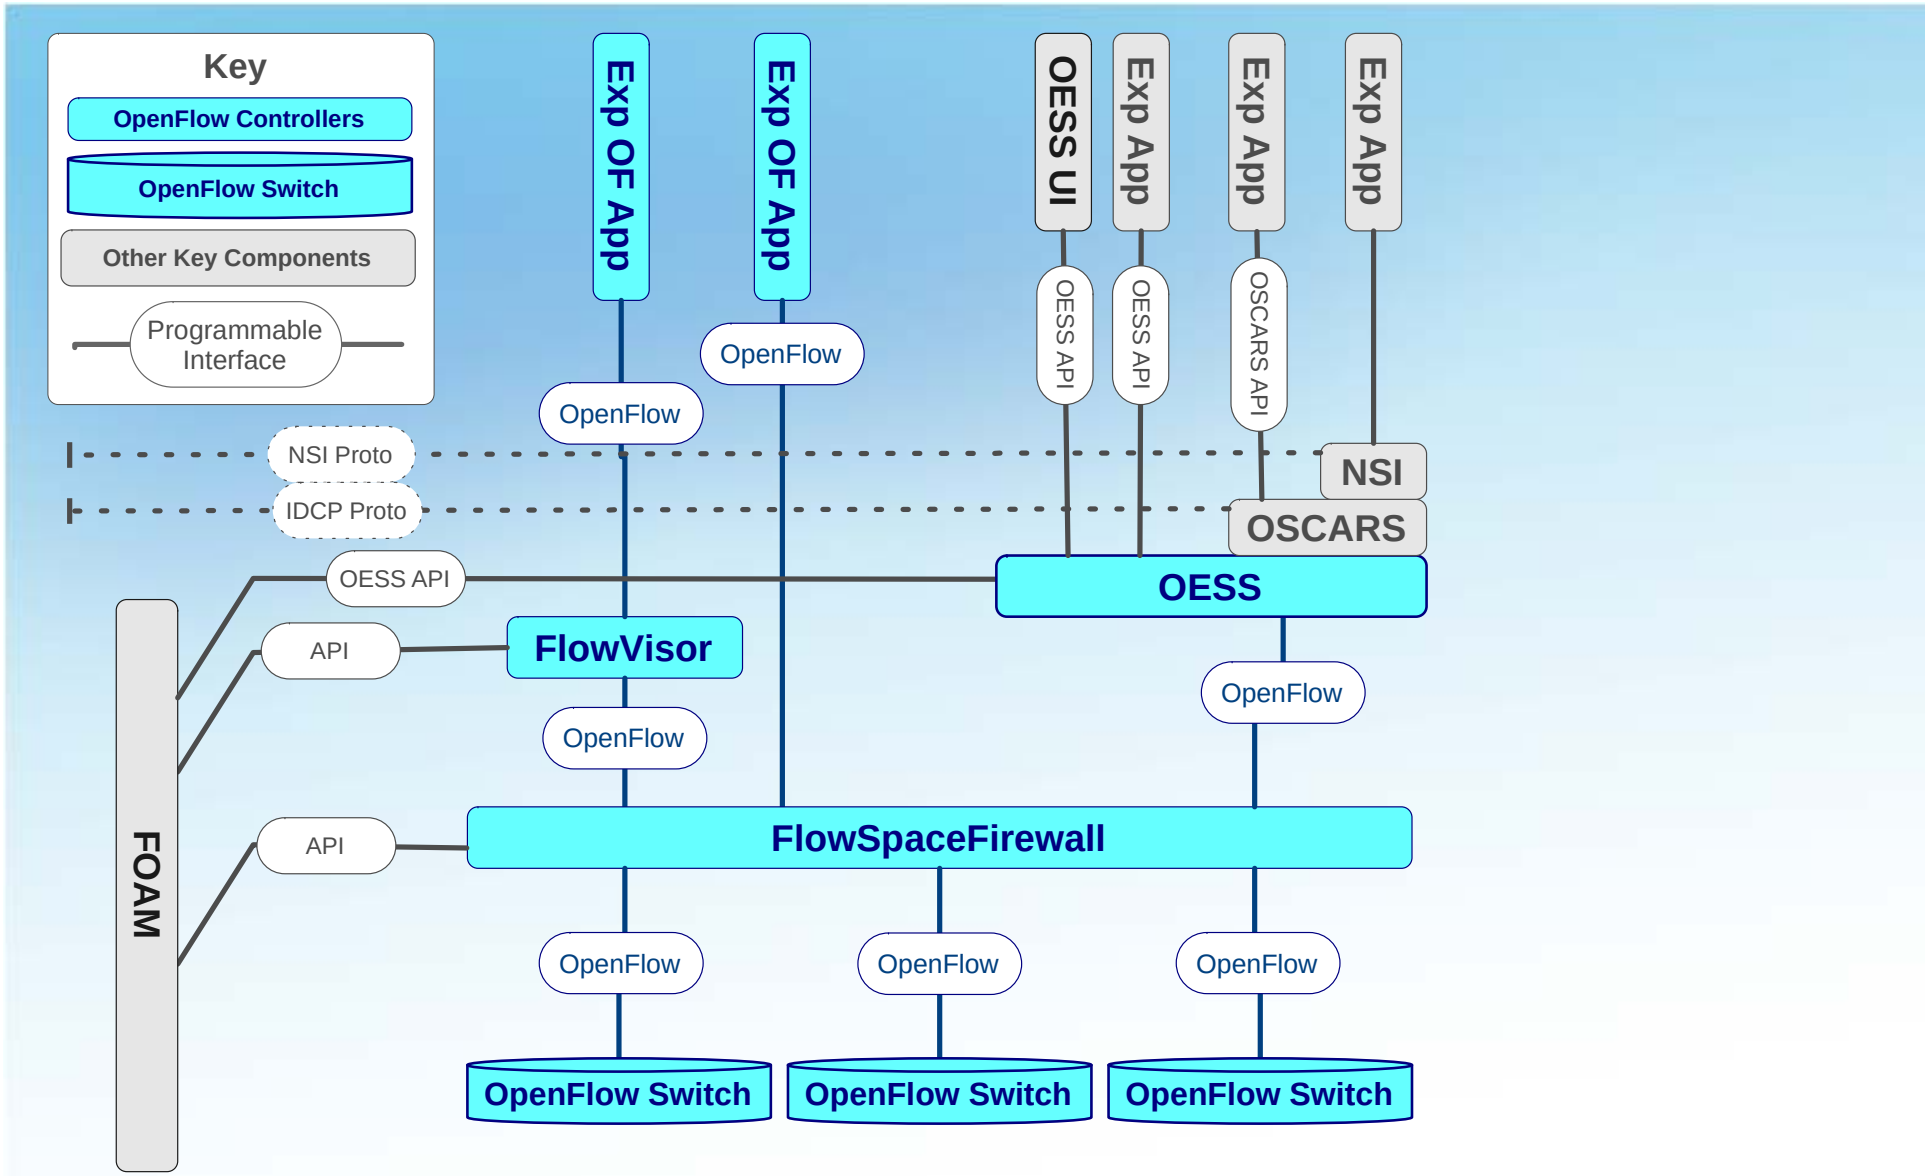

AL2S Software Stack Q1 Late 2014

FlowVisor Performance Issues

- Does not support VLAN Tag range-based policy
  - need 1 policy for every tag on every port in a flowspace
- ~ 1 million policy rules for the AL2S network
  - 28 switches, 10ports each, 4096 policies per port
- Unable to load this many rules in a acceptable time
  - Non-linear

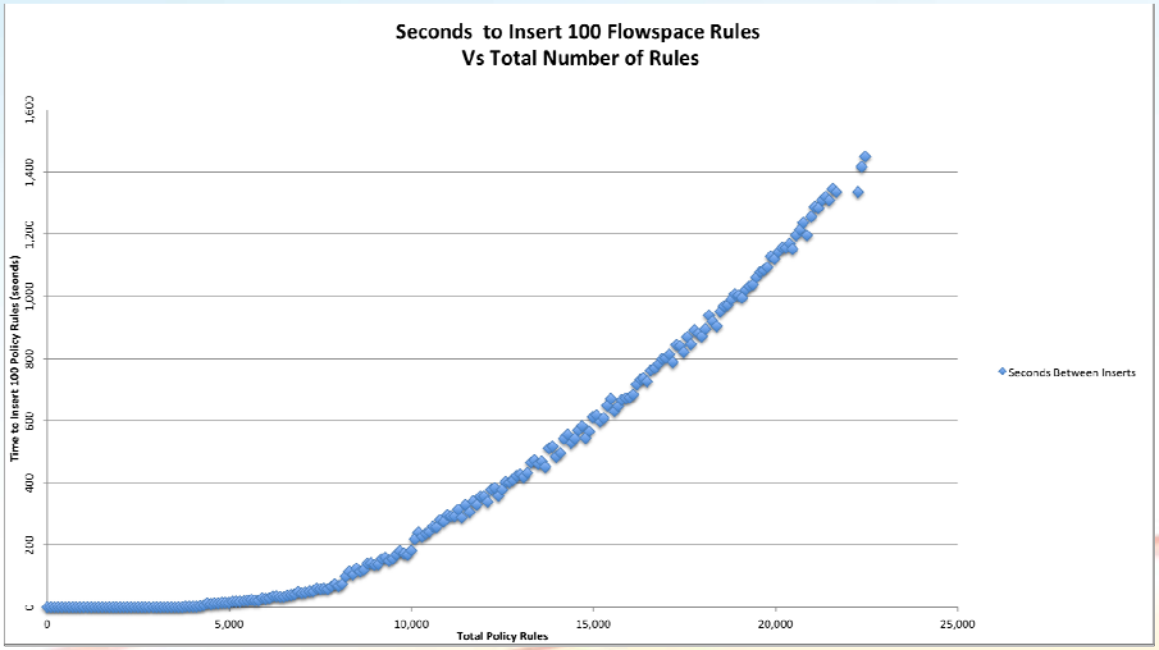

INTERNET 2

FlowVisor Usability Issues

- Policy defined using port numbers not names
  - Port numbers on some systems are ephemeral
  - Difficult for humans to parse
- Policy defined using DPID vs symbolic name
  - DPID on some systems is ephemeral
  - Difficult for humans to parse

```
rule 6182:
FlowEntry[dpid=[00:00:00:a0:a5:7a:d7:34],ruleMatch=[OFMatch[in_port=59590,dl_vlan=4092]],acti
onsList=[Slice:nddi=7],id=[7200],priority=[10],]
rule 6183:
FlowEntry[dpid=[00:00:00:a0:a5:7a:d7:34],ruleMatch=[OFMatch[in_port=59590,dl_vlan=4093]],acti
onsList=[Slice:nddi=7],id=[7201],priority=[10],]
rule 6184:
FlowEntry[dpid=[00:00:00:a0:a5:7a:d7:34],ruleMatch=[OFMatch[in_port=59590,dl_vlan=4094]],acti
onsList=[Slice:nddi=7],id=[7202],priority=[10],]
```

INTERNET 2

Looking beyond FlowVisor

- FlowVisor was designed to provide Flowspace translation
- Translating VLAN tags requires a 1 to 1 mapping
  - Architectural issues behind this
- For AL2S we are more interested in protection than translation
- We need a firewall to keep an OpenFlow application within its defined slice. Slice isolation is essential.
- After working with OnLab, we came to agreement that a separate application would be the most expedient path to resolve
- We need a FlowSpace Firewall.

INTERNET 2

FlowSpace Firewall

- Simple VLAN Tag based flowspace firewall / proxy
- Policy definition and enforcement support range operations
  - < 1,000 policies to support 3 slices using the entire flowspace
- Per slice total rule limits
- Per slice per switch flow modification rate limits *(planned)*
- Built upon FloodLight
- Designed for production use.

Developed by Internet2 with GlobalNOC Software Engineering

INTERNET 2

FlowSpace Firewall Config Example

```
<flowspace_firewall>
   <switch name="foo"  dpid="5" flush_rules_on_connect="false" />
   <switch name="foo1" dpid="2" flush_rules_on_connect="false" />
   <switch name="foo2" dpid="3" flush_rules_on_connect="false" />
   <switch name="foo3" dpid="4" flush_rules_on_connect="false" />

   <slice name="OESS1”>
     <switch name="foo" max_flows="10" flow_rate="1">
       <port name="s5-eth1">
        <range start="1" end="2000"/>
       </port>
       <port name="s5-eth2">
         <range start="1" end="2000" />
       </port>
     </switch>

     <controller ip_address="140.182.45.45" ssl="false" port="6633" />

   </slice>
</flowspace_firewall>
```

Symbolic names reduce policy churn

limits protect network

Range expression for sanity

INTERNET 2

DEMO …

Does this create a platform for innovation?

- Abundant bandwidth to enable innovation? ☑
- Software-defined networking substrate? ☑
- Support data intensive science? ☑
- Virtualization? ☑ In progress
- Integrate network with compute and storage? ☐ Next step

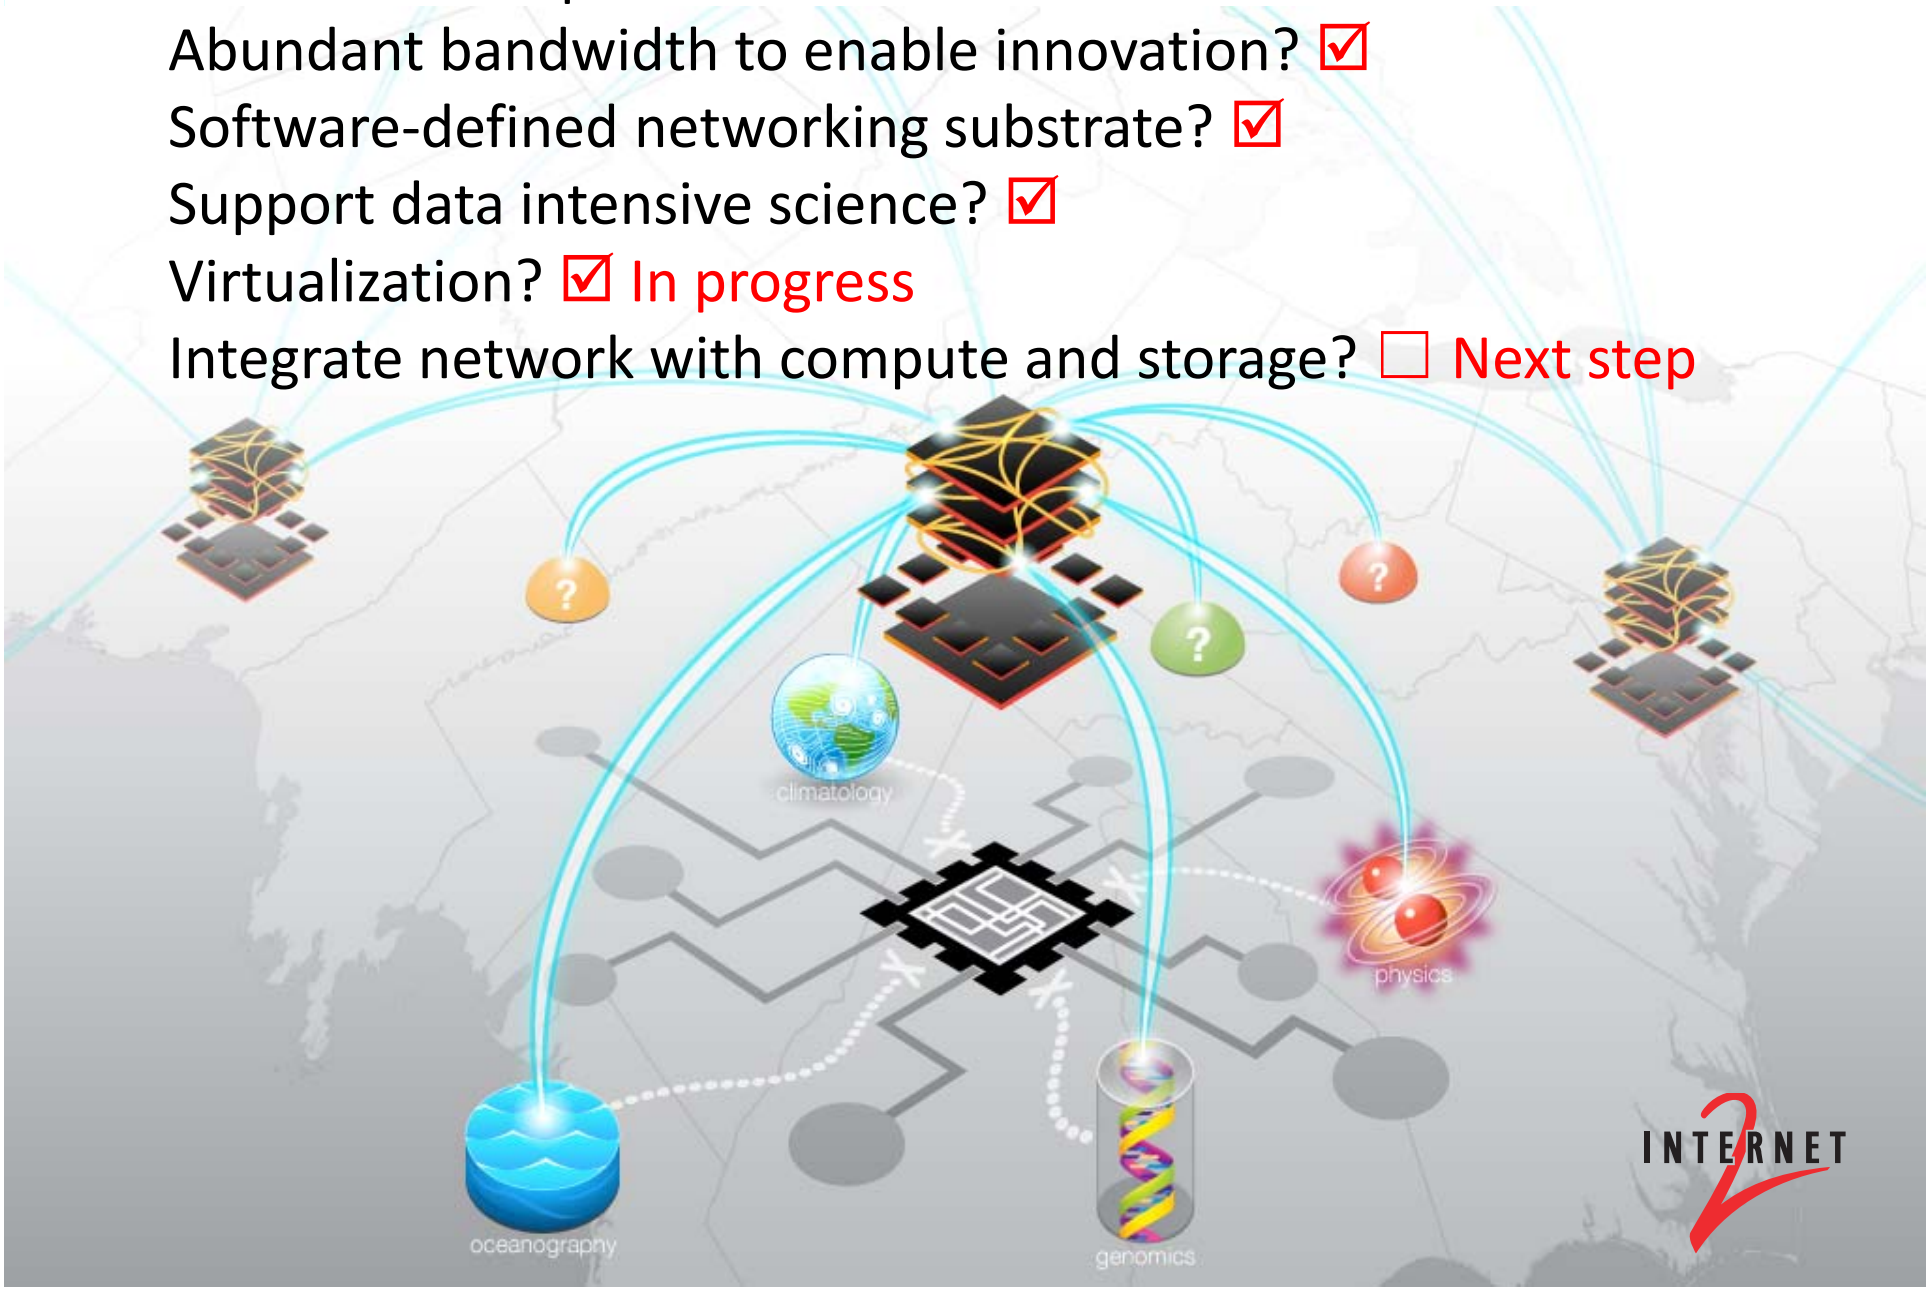

So what does this mean for GLIF Tech?

- We have a great deal of innovation, experimentation, and deployment in areas such as:
  - 100G networking around the globe
  - Software-defined networking
  - Enabling big science flows
- We need understand the various approaches to:
  - Network Virtualization
  - Integration with compute and storage
- Questions to ponder:
  - Does the GLIF community need a common approach to network virtualization?
  - Does the GLIF community need an integrated approach to network virtualization?
  - **More importantly: How does the networking community integrate itself into the compute and storage community?**

INTERNET 2

ENABLING INNOVATION THROUGH NETWORK VIRTUALIZATION
(AND INTEGRATION OF COMPUTE AND STORAGE)

Eric Boyd
Senior Director, Strategic Projects, Internet2
Ed Balas
Manager, Software Engineering, Indiana University

Thank you. For more information,
visit http://www.internet2.edu
or e-mail innovation@internet2.edu

