

---

# GENI Network Stitching Status and Architecture

## GENI GEC16

University of Utah  
Salt Lake City, Utah  
March 19-21, 2013



Tom Lehman  
Xi Yang

University of Maryland  
Mid-Atlantic Crossroads



# GENI Network Stitching Objectives

---

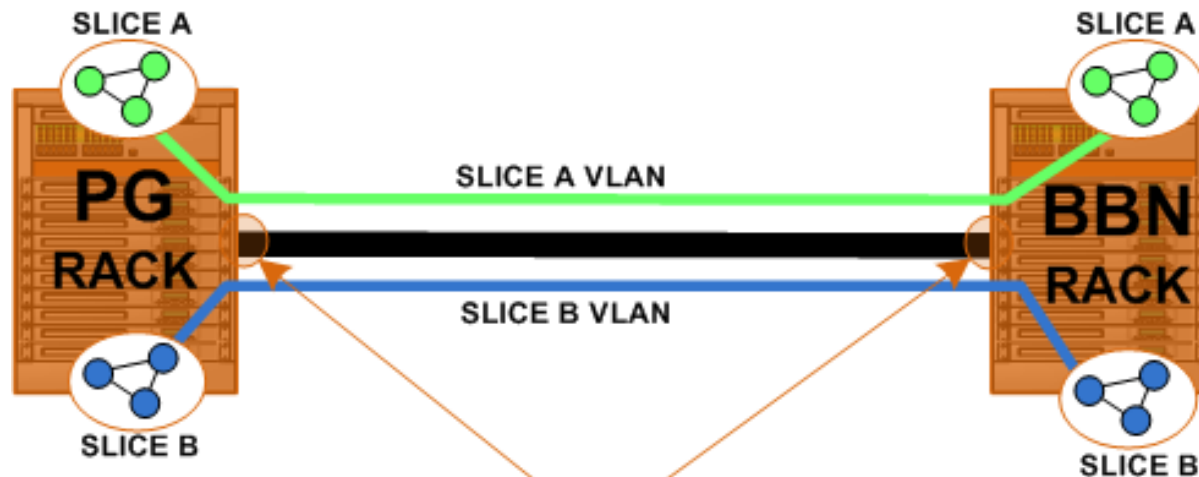
- **Stitching is a GENI Resource issue first, with following objectives from a GENI AM view:**
  - Mapping a "stitching VLAN" to a dynamically create internal slice
  - expressing that "stitching VLAN" on the edge of the GENI Resource (RACK or other Aggregate type)
- **Even if two GENI Racks were directly connected by a dedicated link, stitching would still have to occur**
- **The existence of intermediate networks is a simply a real world factor that the stitching mechanisms must navigate**

# GENI Network Stitching Objectives

## Two directly connected GENI Racks



Still needs stitching to map VLANs to slices

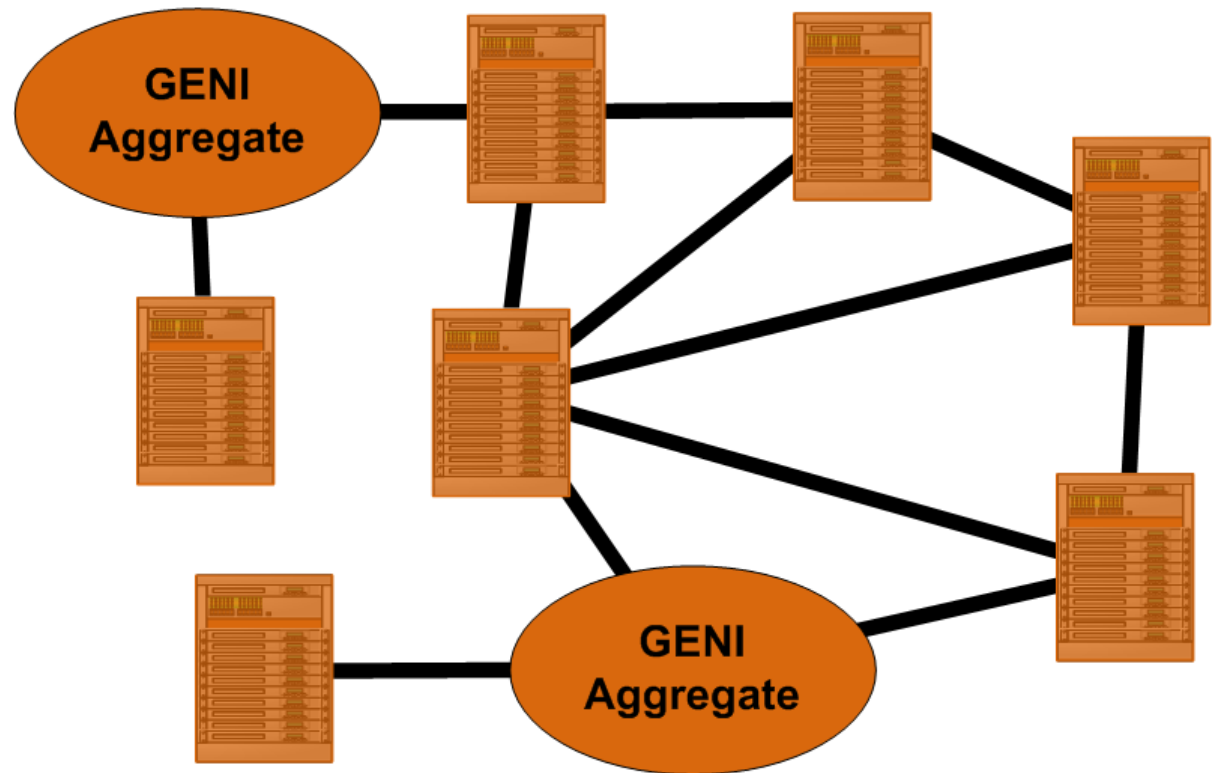


Each Aggregate must express a slice specific VLAN at its edge

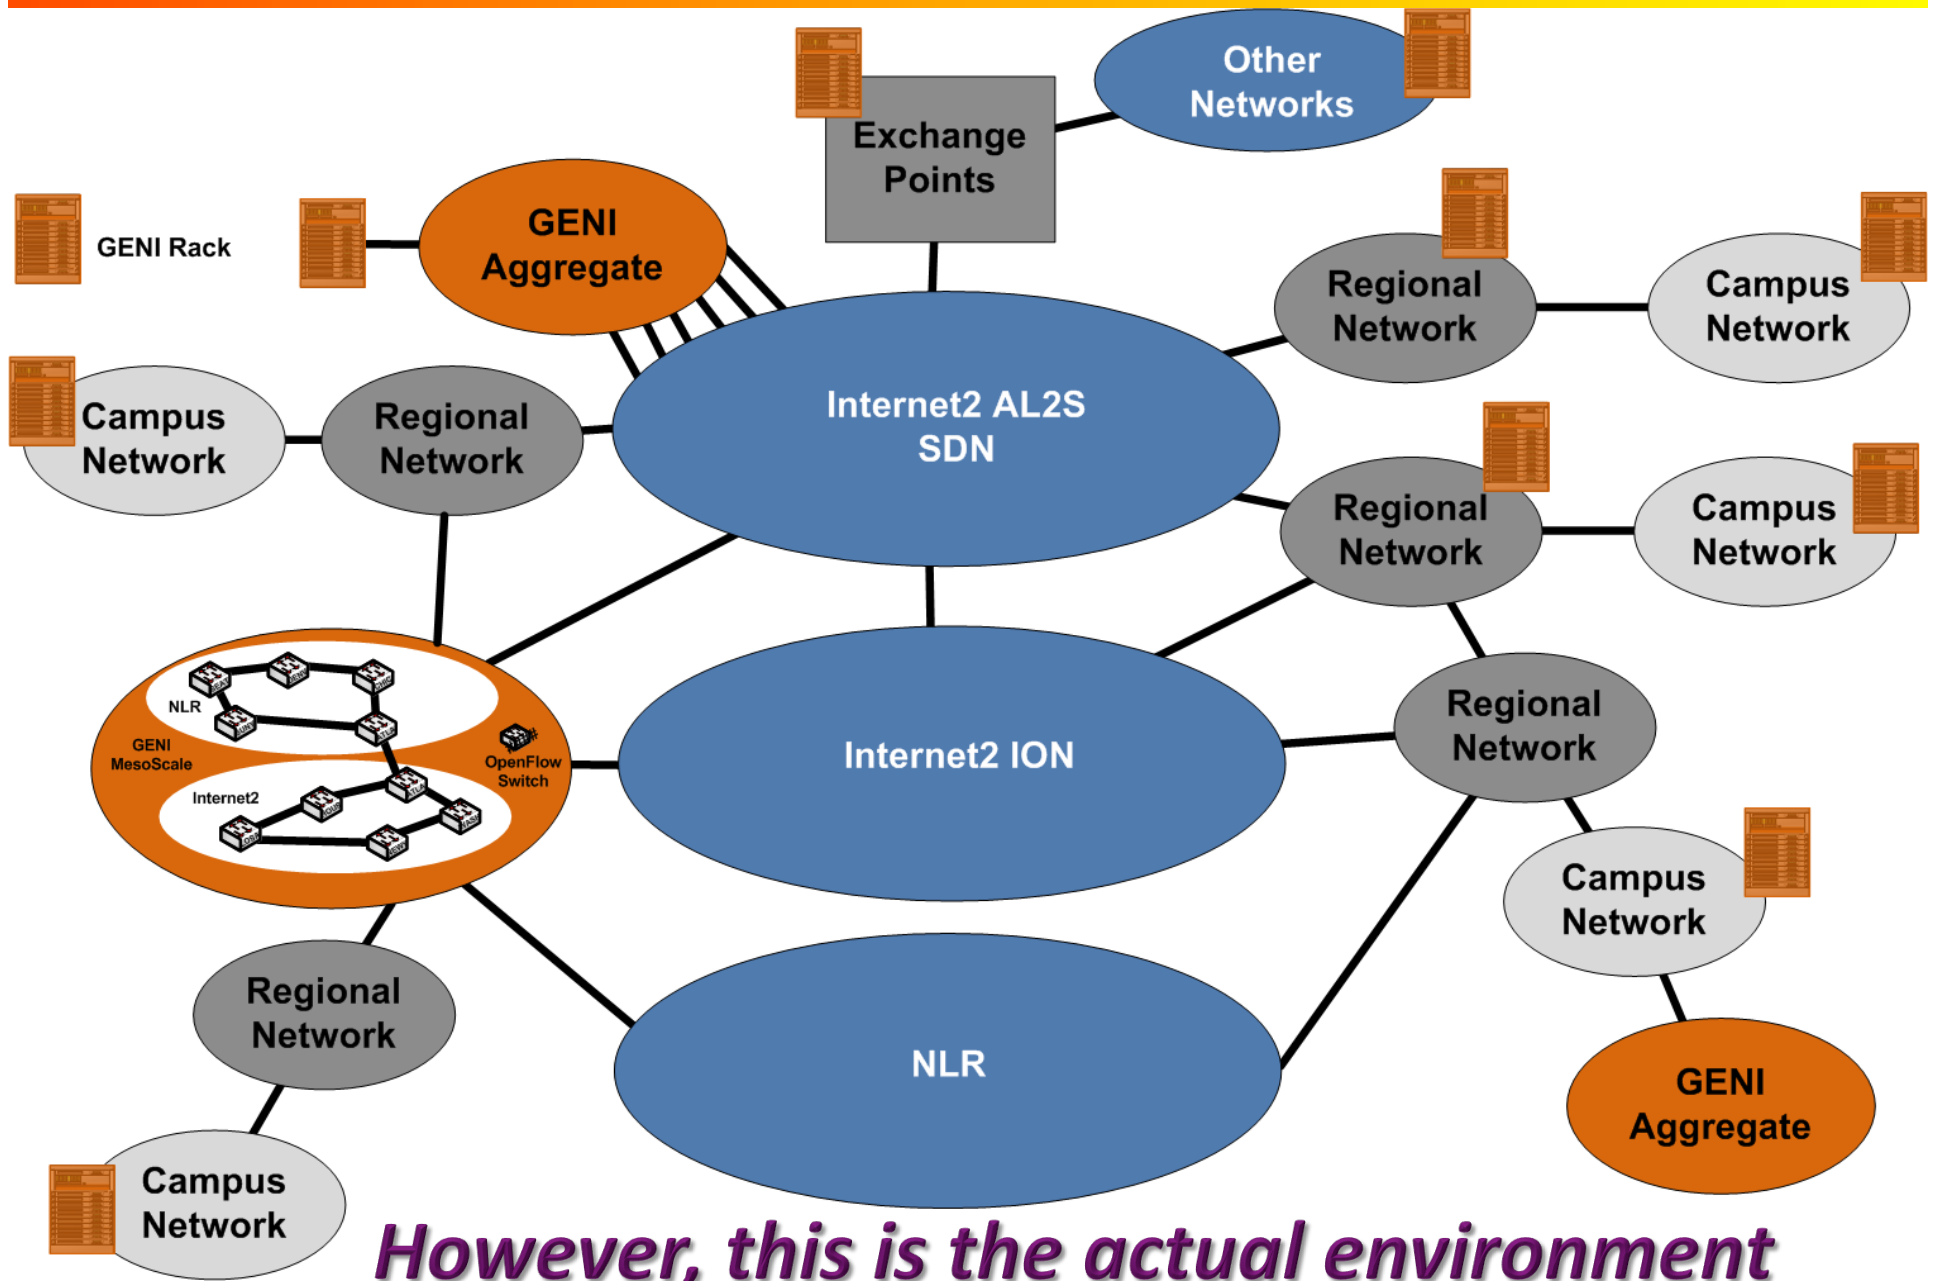
# GENI Network Stitching Objectives

---

- **Stitching for two directly connected GENI Racks would definitely be much simpler**
- **But GENI is not two racks, it is expected to be hundreds of rack and other types of GENI resources**
- **Stitching is required even if there were dedicated links and no intermediate networks**

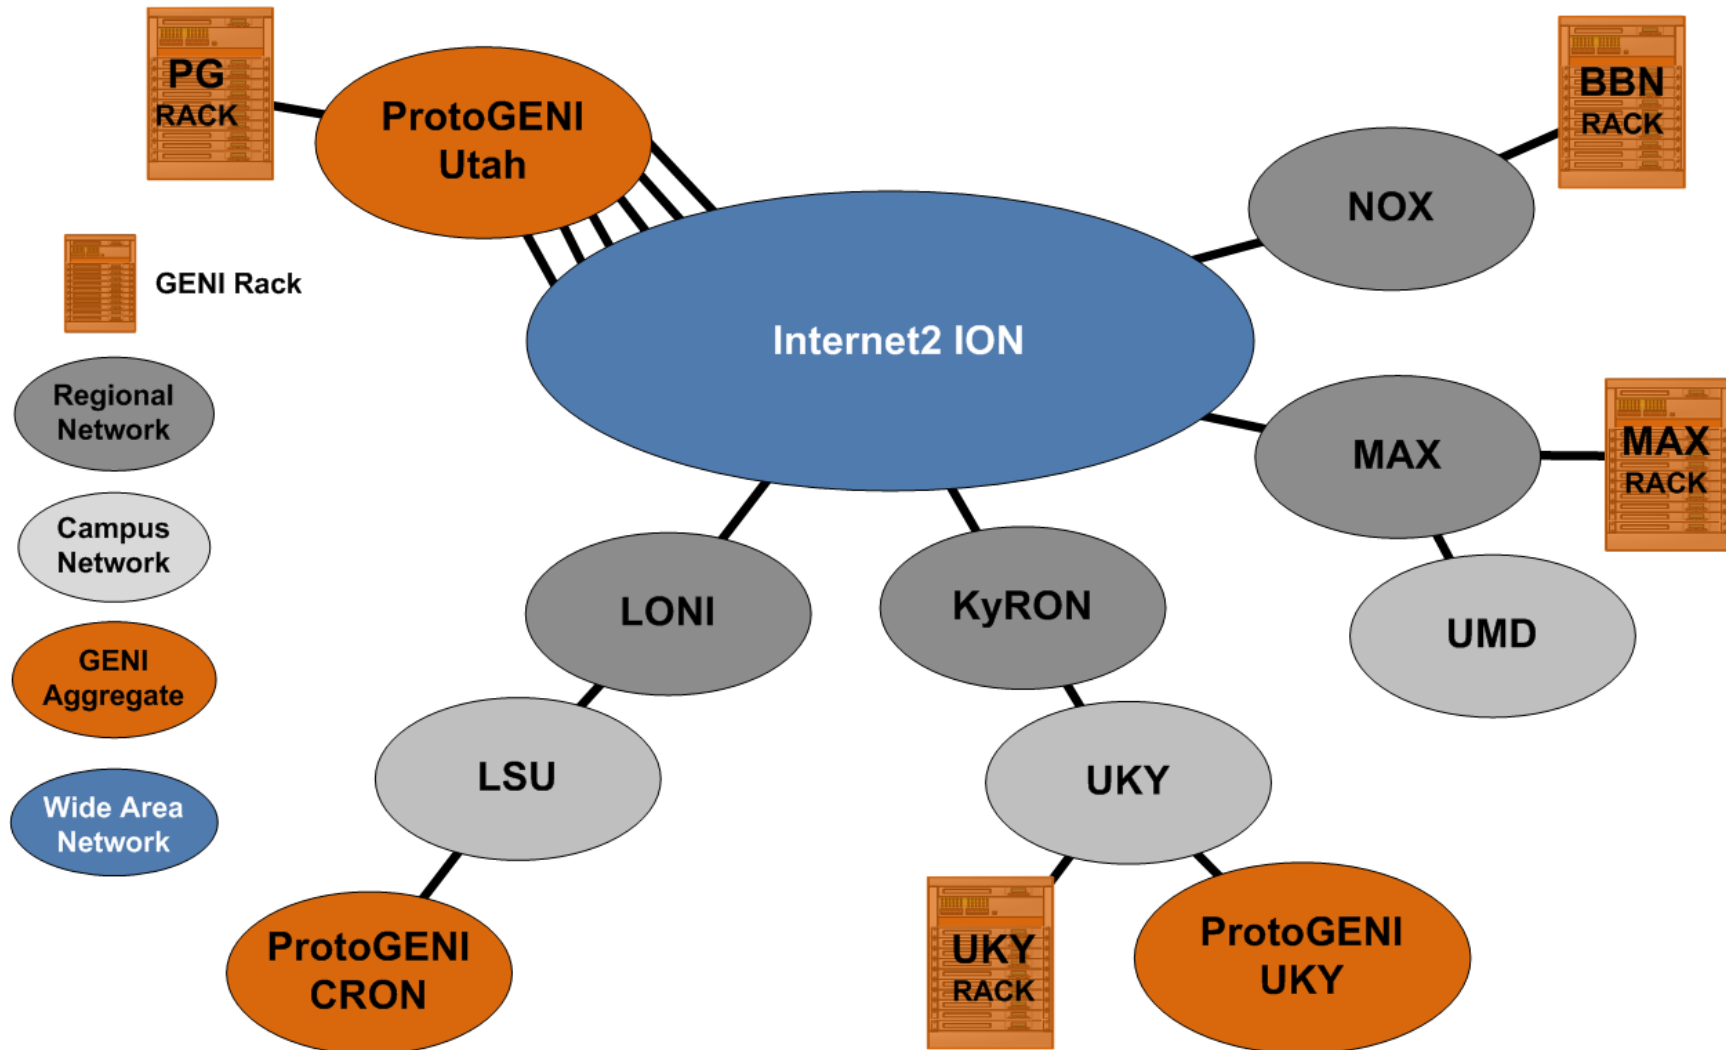


# GENI Network Stitching Environment



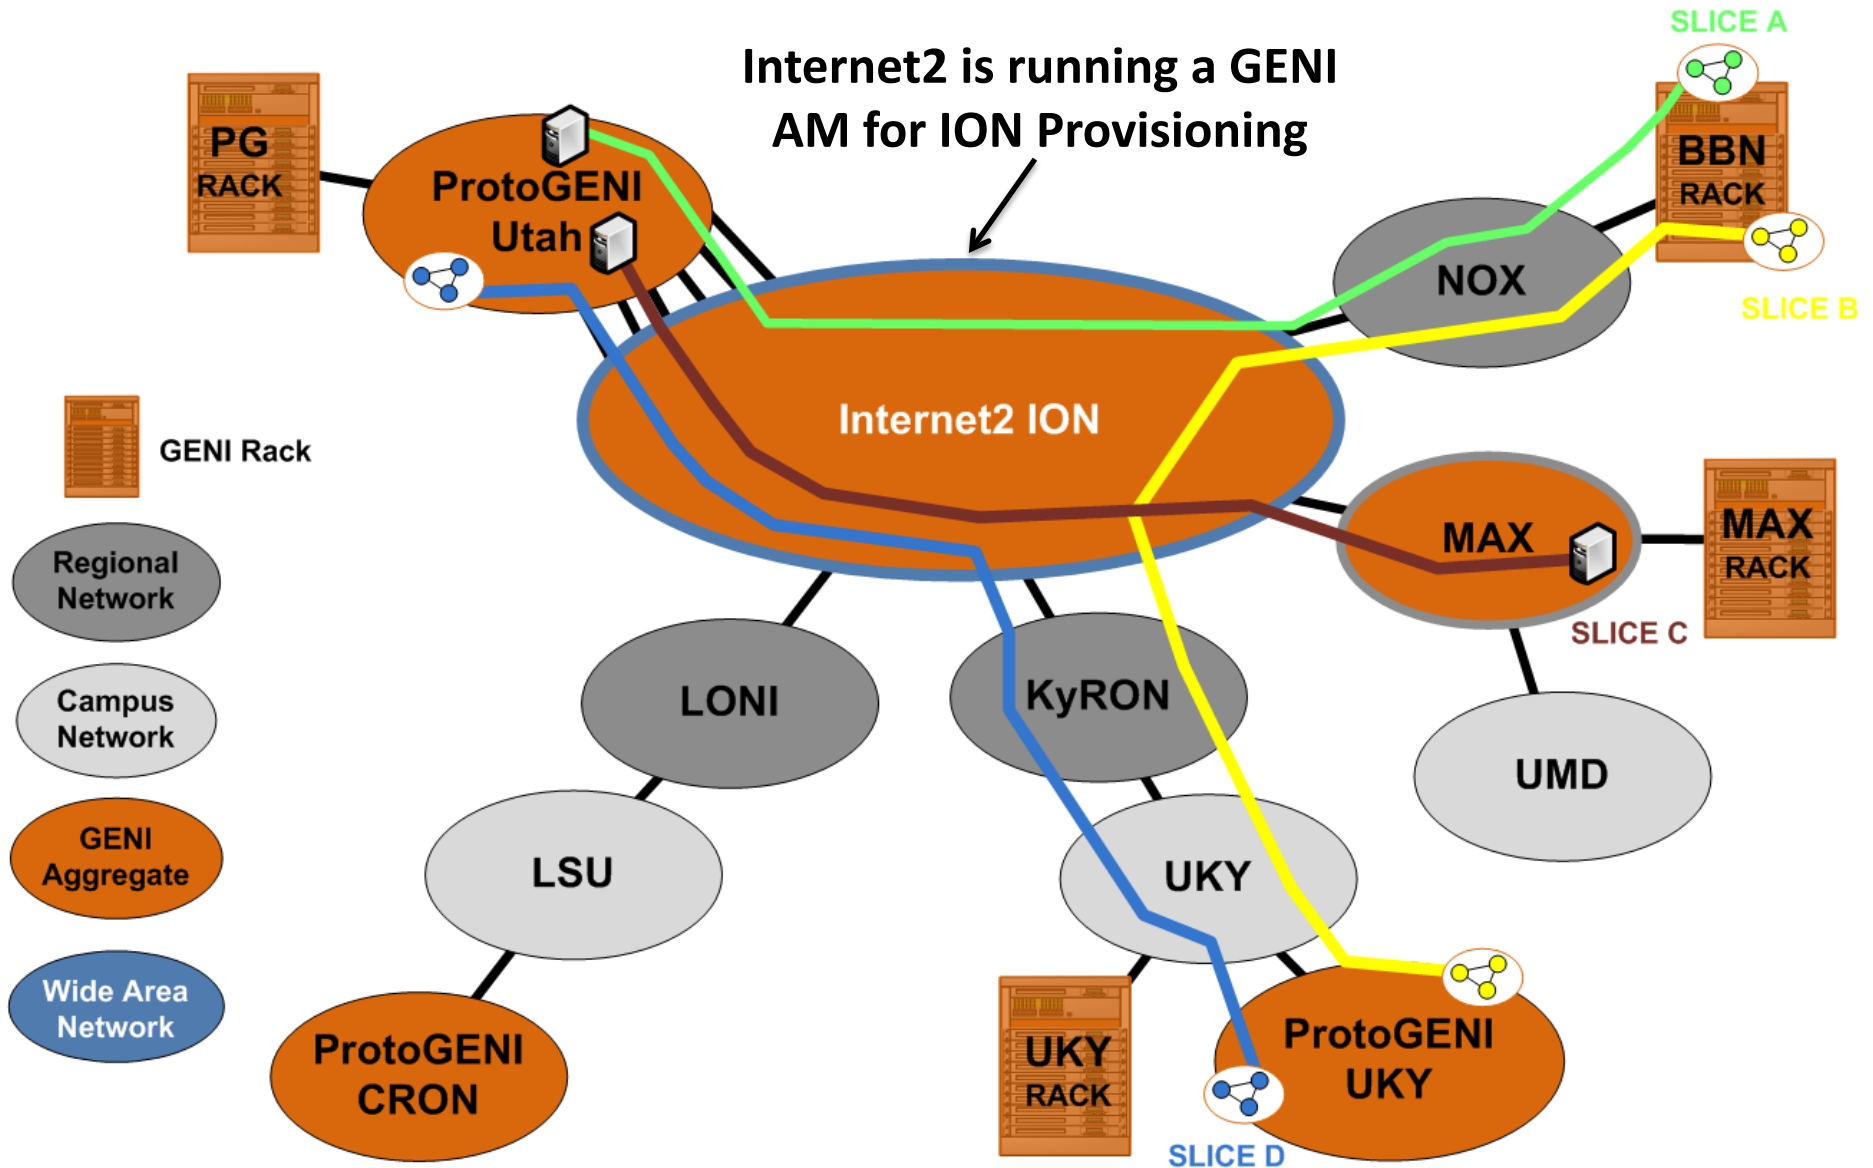
# GENI Network Stitching Environment

## Real World Example



# GENI Network Stitching Environment

## Real World Example – Recent Testing Scenarios



# GENI Network Stitching Environment

---

- **GENI Resources will be dispersed in an organic and Ad Hoc manner across multiple network types: Campus, Regional, Wide Area Networks, and Exchange Points**
- **Multiple methods for provisioning VLANs**
  - static pre-configuration of a range of VLANs
  - dynamic VLAN provisioning (DCN, DYNES, OpenFlow, NSI, others)
  - Campus, regional, wide area, and exchange points will use whatever method they prefer
- **GENI is only API with host, networking, and federation**
  - Key objective of the Stitching Architecture is to be flexible and be able to adapt to multiple and emerging network provisioning technologies to utilize them as part of the GENI control and provisioning framework



# Campus and Regional Networks Stitching Options

---

- **Campus, Regional, Wide Area Networks, and Exchange Points have multiple options to integrate their GENI Resources into stitching**
  - Pre-configure a range of static VLANs
  - Run a Dynamic Network system (DCN/OSCARS, DYNES) with no GENI AM covering
  - Run a Dynamic Network system with a GENI AM covering
  - other methods are also possible
- **In all cases the Advertisement RSpec will reflect how the GENI resource is attached to external resources**
- **The information in the Advertisement RSpec will be used by the stitching infrastructure to facilitate user request and subsequent stitching activities**

# Stitching Extension for RSpecs

---

- **Allows Advertisement RSpecs to define and describe inter-aggregate touch points**
  - Advertisement RSpecs contain mostly static data, e.g., 10G physical inter-connect with vlans 2050-2099 reserved for GENI use
  - Mechanisms to accommodate regions across intermediate networks with static VLAN configurations
- **Allows Request RSpecs and Manifest RSpecs to define slice specific stitching details**
- **Implementation Status**
  - Implemented by ProtoGENI AM, InstaGENI, and MAX AM
  - More work needed with ExoGENI and FOAM AM

# Advertisement RSpecs - Stitching Extensions

---

- Advertisement RSpecs are where the interconnects between GENI Resources and other other network infrastructure is described
- **BBN InstaGENI to ION example (simplified portion):**

BBN InstaGENI Rack Advertisement RSpec main body:

```
<link component_id="urn:publicid:IDN+ion.internet2.edu+link+link-bbninstageni:procurve2:5.24-ion:rtr.newy:ae0:bbn" >  
  <interface_ref component_id="urn:publicid:IDN+ion.internet2.edu+interface+rtr.newy:ae0:bbn"/>  
  <interface_ref component_id="urn:publicid:IDN+instageni.gpolab.bbn.com+interface+procurve2:5.24"/>
```

Stitching Extension has the following:

```
<link id="urn:publicid:IDN+instageni.gpolab.bbn.com+interface+procurve2:5.24">  
  <remoteLinkId>urn:publicid:IDN+ion.internet2.edu+interface+rtr.newy:ae0:bbn</remoteLinkId>  
  <vlanRangeAvailability>3747-3749</vlanRangeAvailability>  
  <vlanTranslation>>false</vlanTranslation>
```

- **Regions of statically mapped VLAN ranges will be reflected in the Advertisement RSpec**
  - as shown above for VLAN range 3747-3749 mapped thru NOX

# GENI Network Stitching Responsibilities

---

- **GENI Rack/Aggregate Manager has the following responsibility with regards to stitching:**
  - map/attach the "stitching VLAN" to the slice specific resources inside the GENI Rack
  - express the "stitching VLAN" at the edge of the GENI Rack (facing the locally attached network)
  - nothing extra required from GENI AM using APIv2
- **For GENI AM APIv3, there will be some additional requirements for the GENI AM to support negotiation:**
  - specific responses to allocate and reserve messages to facilitate the GENI Stitching Infrastructure finding a slice specific "stitching vlan" value that works across multiple GENI aggregates/racks
  - this will be more critical as the number of GENI Racks and simultaneous experimenter slices increase

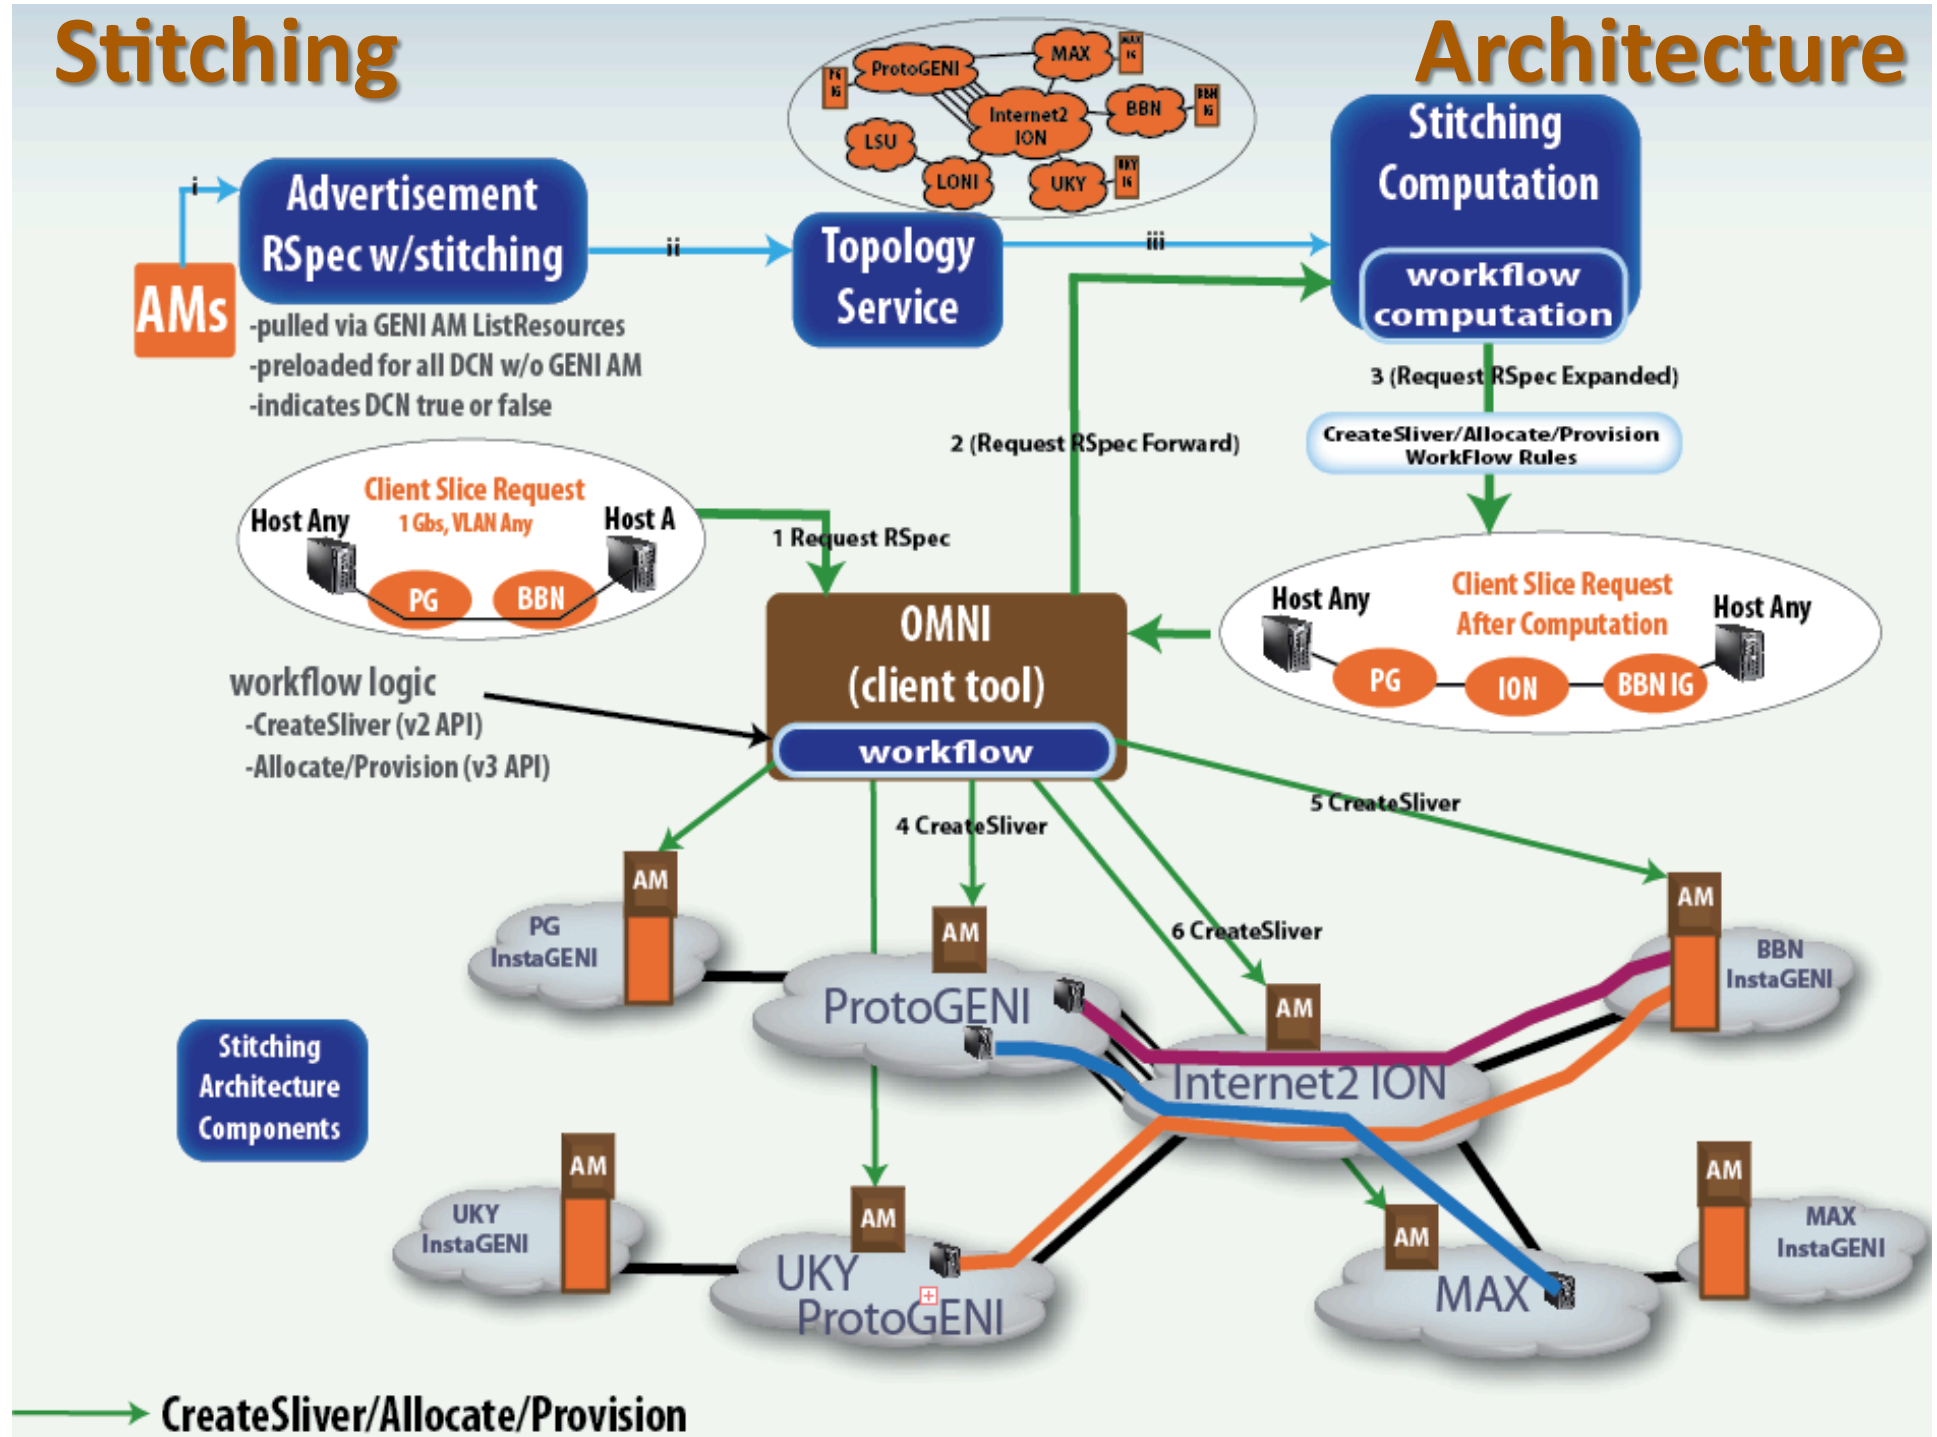
# GENI Network Stitching Responsibilities

---

- **The Stitching Infrastructure has the following responsibilities with regards to stitching:**
  - Find a "stitching VLAN" value that will work as part of dynamic slice instantiation
  - Interact with the GENI AM via standard AM API
  - operate in a world where the campus, regional, wide area networks, and exchange points use various methods to get "stitching VLANs" provisioned
- **Stitching Infrastructure in this context is: Stitching Extensions in RSpecs, Topology Services, Stitching Computation Services, Stitching Library for GENI Clients**
  - objective is to be transparent to user, they can just supply what they want, stitching infrastructure will figure out the details

# Stitching

# Architecture



# Stitching Architecture Components

---

- **Stitching Extension**
  - for Advertisement, Request, and Manifest RSpecs
- **Stitching Computation Service**
  - with workflow rule generation
- **Workflow Execution Stitching Library**
  - for integration into clients such as OMNI, Flack, GENI Portal, and others
- **Stitching Topology Service**
  - collects Advertisement RSpecs from GENI AMs
- **GENI AM API (v2, v3) Stitching Processing**
  - v3 has features to facilitate Stitching with negotiation

# Stitching Extension for RSpecs

---

- **Stitching Extensions defined (v0.1)**
  - used for Advertisement, Request, and Manifest RSpec
  - Implemented by ProtoGENI, InstaGENI, MAX
  - More work needed with ExoGENI and FOAM
- **Stitching Extensions v2 under review**
  - adds shared vlan/openflow vlan support
  - adds optional capability elements to indicate things such as vlan producer/consumer
- **A few open items relating to stitching extensions need discussion. These are listed on slides at end of this presentation**



# Stitching Computation Service (SCS)

## Stitching Computation Service



1 Request RSpec



Client tool logic required to execute workflow dependency and interact with Aggregate Managers

2 (Request RSpec Forward)

SCS Computation Request  
*slice urn*  
*request rspec*  
*hop exclusion list (optional)*  
*hop inclusion list (optional)*

note: exclusion/inclusion allows client to indicate preferences at the aggregate, node, port, or vlan level



3 (Request RSpec Expanded)



SCS Computation Return

*request rspec expanded*  
*workflow data (optional)*

-dependency set for each stitching path in the request rspec  
-considers the following aggregate and element capabilities when developing dependencies:

--vlan producer/consumer, vlan translation in the context of specific aggregates and elements, vlan translation in the context of topology of slice

-evaluates for loop free result

# Stitching Computation Service

---

- **Uses Advertisement RSpec information to perform slice instantiation specific computations**
  - manually loaded due to lack of topology service
  - can also include DCN topologies
- **Expands client provided RSpecs to include intermediate GENI Aggregates and DCN networks required for stitching**
  - adds stitching extension automatically even if client does not provide
- **include and exclude request options at aggregate, node, port, vlan levels**
- **XML-RPC API**

# Stitching Computation Service

---

- Includes a workflow data generation option, which provides guidance to client tools (Omni) for interaction with AMs.
- Workflow data returns array of recursive dependencies indicating proper sequence and associated actions for Aggregate requests
- **The following information/capabilities are utilized to calculate dependencies**
  - vlan producer/consumer status
  - vlan translation capabilities
  - vlan continuity (position of specific aggregates in the expanded RSpec and available vlan range size)
  - evaluates for loop free result
- **Advertisement RSpec information is utilized as needed to calculate expanded RSpec and dependencies**
- **Implementation Status**
  - prototype under test now
  - details available here:
    - <http://geni.maxgigapop.net/twiki/bin/view/GENI/NetworkStitchingAPI>

# Workflow Stitching Library

---

- OMNI developers have generated a stitching library which can process SCS Workflow data and submit proper requests to Aggregates
- Similar functionality can be integrated into other client tools (such as Flack and GENI Portal) once OMNI integration is complete
- **Implementation Status**
  - prototype under test now

# Topology Service

---

- **Stitching Topology Service**

- objective is to collect Advertisement RSpecs from all GENI AMs
- this will allow the building of larger views (global if desired) of inter-connected GENI AM space
- also plan to integrate topology data collection and views across the dynamic networking/perfsonar space to enable GENI stitching to better leverage these capabilities

- **Implementation Status**

- refining requirements and evaluating options for topology service
- some initial implementation work done by GEMINI project to collect AM Advertisement RSpecs
- will be a focus after GEC16

# GENI AM APIv3 – Stitching Negotiation

---

- **GENI AM APIv3**

- changes AM interaction from "CreateSliver" to "Allocate, Provision" paradigm (similar to the GetTicket, RedeemTicket model)
- when trying to find compatible resources (in this case bandwidth and vlans) across multiple independent aggregates, there is value in being able to "negotiate"
- we plan to incorporate this "negotiation" feature into future stitching operations.
- Current Stitching implementations do not require this negotiation feature because we can intelligently utilize the various aggregate capabilities such as VLAN translation when sequencing Sliver requests.

# Existing Dynamic Network Services/ Infrastructure

---

- **Dynamic Circuit Network (DCN) infrastructure available today**
  - uses Inter Domain Controller Protocol (IDCP)
  - Implementation based on OSCARS and DRAGON software.
  - Participating networks include Internet2 Advanced Layer2 Network Services (AL2S), DYNES project deployments (approximately 40 campus, regional networks, exchange points), DOE ESnet, multiple international networks (RNP, JGN2), and others
- **Other Dynamic Networks Services are expected in the future**
  - additional deployments of DCN based services at Exchanges such as WIX and MANLAN
  - systems based on other technologies such as OpenFlow, Network Service Interface (NSI), others
- **Key objective for GENI Stitching Architecture is to be able to use existing and future Dynamic Network Services**

# Existing Dynamic Network Services/ Infrastructure

---

- **Internet2's Advanced Layer2 Network Services that we are leveraging**
  - ION Service
  - SDN/OpenFlow Network (AL2S)
- **Internet2 is running a GENI AM for ION on a trial basis to facilitate GENI Stitching**
  - allows for GENI stitching across ION and SDN/OpenFlow Network
  - uses the OSCARS services underneath the GENI AM
  - Only point-to-point EVPL (Ethernet Virtual Private Line) services available
- **The Internet2 SDN/OpenFlow AL2S network has the potential to provide additional services (i.e. multipoint)**



# GENI Networking Stitching Schedule

---

- **Stitching being tested now between GENI Racks and other Aggregate end points**
- **Planning for a beta release of OMNI with stitching after GEC16 for prototype testing**
- **Hope to make stitching services available for general use before GEC17**
  - limited set of participating GENI Aggregates and intermediate networks will be configure for initial release and testing
  - This will include initial implementations for AM RSpec Stitching Extension support, Stitching Computation Service, Omni Stitching Library
- **Stitching Topology Service implementation to follow**

# GENI Networking Stitching Schedule

---

- **After GEC16**
  - AMs using GENI AM APIv3 for negotiation
  - GENI Portal and Flack stitching integration
  - Stitching with ExoGENI and FOAM Aggregates
  - Enhanced implementations of Stitching RSpecs and AM support, Computation Service, Topology Service, Stitching Library

# Stitching Open Items

---

- **InstaGENI Advertisement RSpecs handling of logical interfaces and associated mapping to stitching extension**
  - InstaGENI Ad RSpec currently advertises duplicate links in some situations
- **ExoGENI and FOAM support for Stitching Extension in Advertisement RSpec, Request RSpec, and Manifest RSpec**
  - need to discuss how main body of RSpec will be mapped to stitching extension
- **Document all the rules for generation of stitching extension from the RSpec main body**
  - the construction of "free form" fields for the interface elements in the Advertisement RSpec, for the inter-aggregate interfaces
  - implied internal connectivity based on use of wildcards in interface urns
  - Confirm that all RSpecs should have info from all Aggregates, in both main body and stitching extension for multi-aggregate slices
- **Define stitching changes as related to GENI AM APIv3**
- **FOAM/Openflow AM which supports stitching and VLAN level slice provisioning**

# Stitching Open Items

---

- **Define best practices for configuration and use of multi-aggregate layer2 connections**
  - do we want to formalize any best practices as it relates to spanning tree or BPDU filtering?
  - rules for termination/extension of a stitched vlan beyond the terminating GENI resource?
- **Add support for Dynamic Networks (OSCAR/PerfSonar based) which do not run a covering GENI AM**
  - A GENI Advertisement RSpec will be auto generated from the PerfSonar Topology Service
  - Special workflow rules will be used to allow OMNI to initiate provisioning from a GENI AM which spans multiple DCN domains
    - changes to SCS and OMNI Stitching library required

---

**Thank you**

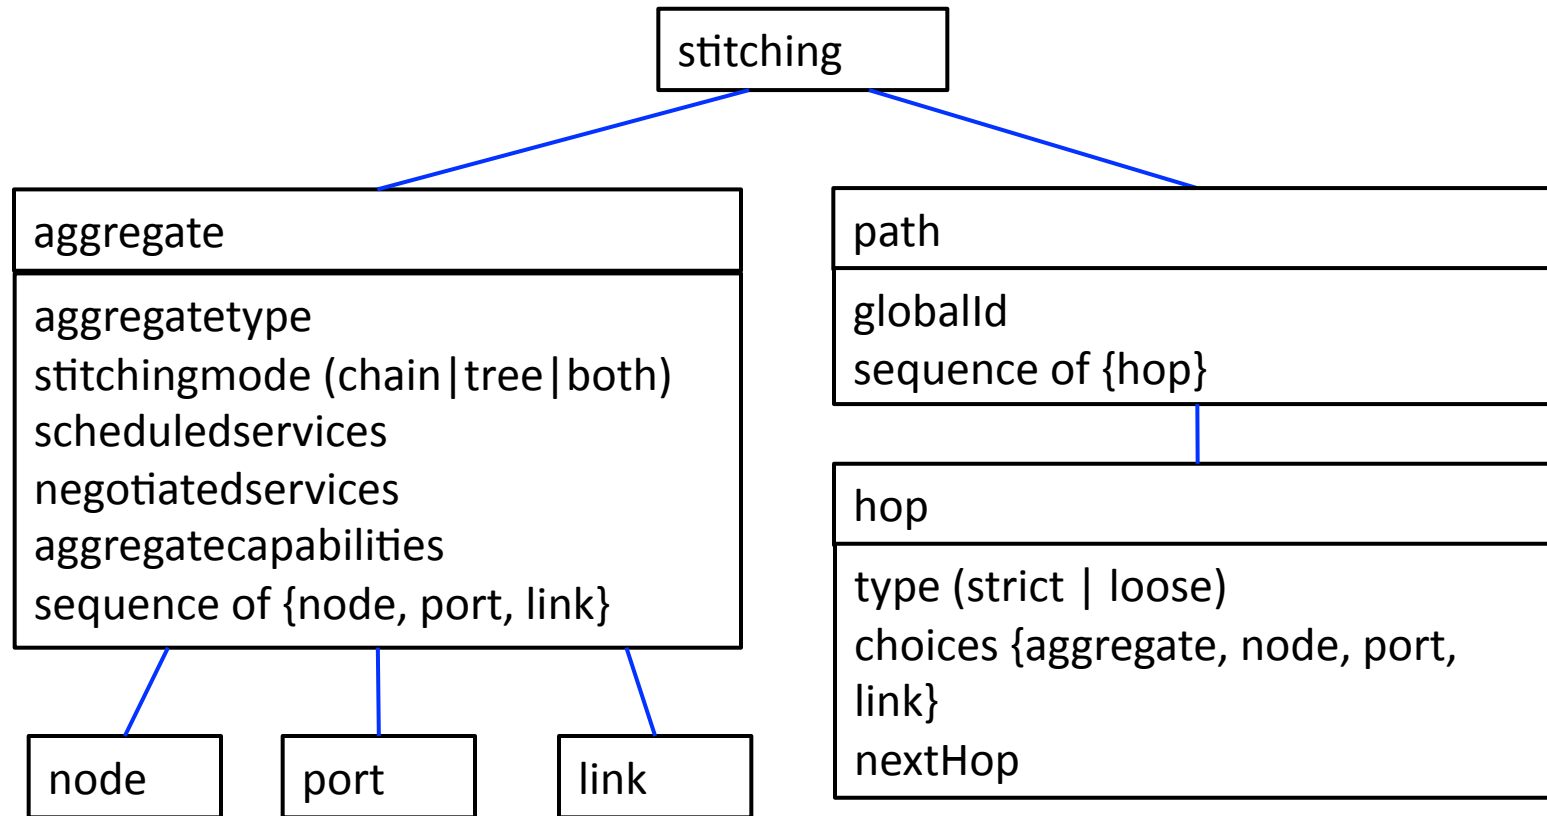
---

# GENI AM API Overview

---

- **Resource Specification (RSpec)**
  - Advertisement RSpec: used by Aggregate Manager (AM) to publish their resources
  - Request RSpec: used by client tools to describe requested experiment resources
  - Manifest RSpecs: used by AM to describe instantiated experiment topology
  - Common XML Schemas used by all GENI AMs
- **GENI AM API messages**
  - messages for credentials and RSpec exchange (advertisement, request, and manifest)

# Stitching Extension Schema



*technology specific details here*

<http://www.geni.net/resources/rspec/ext/stitch/2/stitch-schema.xsd>

# ION Aggregate Manager

- Based on MAX AM which uses the SFA PlanetLab Framework

