GENI Stitching
Status for Developers
GEC 17
Madison, Wisconsin

Aaron Helsinger
July 22, 2013
www.geni.net

The Demo was Impressive Because..

GENI Stitching is

- Fast
- Instructions are posted so you can do it
- Connect InstaGENI to ExoGENI
- Connect a non GENI host to a GENI VM
  - Across CENIC to a machine at Stanford

Demo Details

- Stitching to Fixed Endpoint Works
- Stitching to ExoGENI Works
- Stitcher has some improvements coming
- This implements the GENI Stitching Architecture

Stitching since GEC16

- Released Omni 2.3.3 with stitcher
- Operational Testing
  - Ongoing
  - http://groups.geni.net/geni/wiki/GeniNetworkStitchingTestStatus
- Support for ExoGENI
- Support for fixed endpoints
- Stitcher Fixes and Improvements
  - Options to control timing of ION aggregate status checks & retry
  - Better aggregate error handling
  - Cache slice credential

ION Aggregate Manager

- ION is an Internet2 service that uses OSCARS to do dynamic circuits and does VLAN translation
- Tom Lehman and Xi Yang of MAX wrote an SFA based aggregate manager that translates GENI calls to OSCARS calls
- Internet2 operates this aggregate
  - *Other OSCARS networks (regionals?) could use it too*
- This is powerful:
  - GENI experimenters can connect arbitrary GENI ION endpoints when they need them.
  - This enables stitching to non GENI resources (i.e. through CENIC to a node in Stanford as in the demo).

Thank you MAX and Internet2!

Experimenter View: Creating a Circuit

2. Send Path Request to Stitching Computation Service (SCS)

1. Simple Request

```
<link client_id="mylink">
  <component_manager…
  <component_manager…
```

SCS

3. Get Expanded Request

Tool

Aggregate 1

4. Send Request to Aggregate 1

5. Get Manifest

Aggregate 2

6. Repeat for Other Aggregates

```
<link client_id="mylink">
  <component_manager…
  <component_manager…
…
<stitching>
  <path id="mylink">
    <hop id="1">
      <link id="switch1:port1"
…
        <vlan…>3747</vlan…>
```

7. Manifest Back

Finding a Path

- Tool computes possible paths and VLANs
- Order of reservations impacts chance of success
  - Some aggregates must pick the VLAN
  - VLAN translating aggregates can go last
- Your tool does all this for you

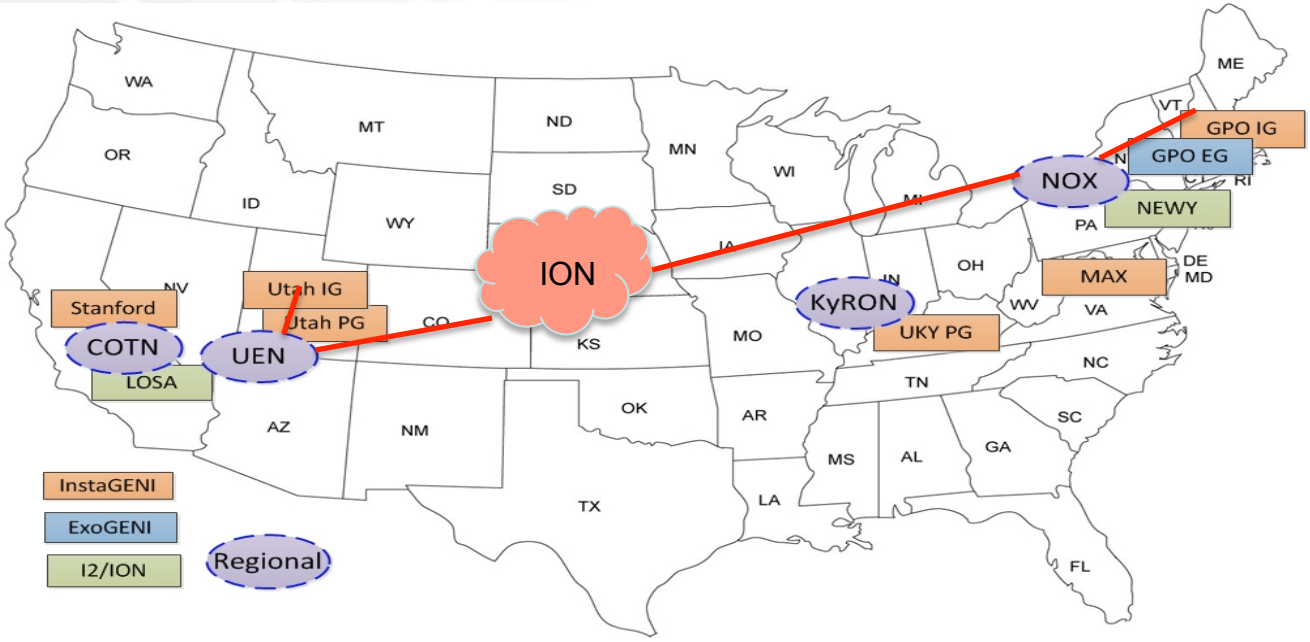

Stitching Computation Service

- Stitching Computation Service (SCS) for path and workflow computation
  - Includes many heuristics to optimize path, chance of success
  - Allows excluding particular connection points, VLANs
- Tom Lehman and Xi Yang wrote this optional service
  - Other tools may use different heuristics
- Stitcher uses the SCS

  http://geni.maxgigapop.net/twiki/bin/view/GENI/NetworkStitchingAPI

ION AM and SCS since GEC16

- ION AM
  - Full IDC/OSCARS v6 support, including `ErrorReport`
  - Policy Enforcement based on Ad Rspec
    - Requested endpoints, bandwidth, VLANs must match Ad
  - `RenewSliver` implemented
  - Released GENI ION-AM v1.0-alpha
- SCS
  - ExoGENI Support
  - `GetVersion` added to API
  - Released GENI SCS v1.0-alpha
  - http://geni.maxgigapop.net/twiki/bin/view/GENI/NetworkStitchingSoftware

ExoGENI Support

- ExoGENI is different
  - URN structure is different
  - Each AM has 2 URNs
  - Each AM has 2 URLs
  - Manifest has different hops than request
    - Adds an internal hop
    - Drops non local hops
    - Ignores 2 apparently local hops
  - Hop IDs are not preserved
    - But hop->link IDs are
- We have stitched to EG GPO.
  - Any EG site should now work

Fixed Endpoints

- Given a switch/port/VLAN, we can stitch there
- Must add any new ION endpoints to OSCARS, ION AM, SCS
- Must add the new remote/fixed endpoint to SCS
- Stitcher: added a new option to make InstaGENI aggregates happy with a non-GENI endpoint

Managing Networks

3 Options for controlling circuits at networks:

- Run an Aggregate over Dynamic Circuits
  - We have an OSCARS aggregate. Others possible.
- Run an Aggregate for Static Circuits
  - Network manages the VLANs
- Delegate VLANs to the peer GENI aggregates to manage
  - Set up the VLANs, configure them at the aggregates

VLAN Delegation

- Networks with no aggregate themselves contribute their circuits, VLANs
- Record details for aggregate configuration
- Experimenters might want to know what networks their circuit traverses
- These networks might want to be able to see their contribution, track how GENI is using the circuits
- Info goes in a config file, and then to Ad RSpecs, wiki

VLAN Delegation Config File

- VLAN Delegation / Intermediate Networks
  - XML Schema to represent information
  - Script fills in config file from Ad RSpecs as needed

```
<link id="NOX1">
  <interfaceA>
    <componentId>urn:publicid:IDN+instageni.gpolab.bbn.com+interface+procurve2:5.24</componentId>
    <aggregateManagerUrl>https://www.instageni.gpolab.bbn.com:12369/protogeni/xmlrpc/am/2.0</aggregateManagerUrl>
    <description>BBN GPO-IG to Internet2 ION via NOX</description>
  </interfaceA>
  <interfaceZ>
    <componentId>urn:publicid:IDN+ion.internet2.edu+interface+rtr.newy:ae0:bbn</componentId>
    <aggregateManagerUrl>http://geni-am.net.internet2.edu:12346</aggregateManagerUrl>
    <description>Internet2 ION to BBN GPO-IG via NOX</description>
  </interfaceZ>
  <vlanProviderAtoZ id="urn:publicid:IDN+nox.org"/>
  <capacity>1000Mbps</capacity>
  <maximumReservableCapacity>1000Mbps</maximumReservableCapacity>
  <minimumReservableCapacity>1Mbps</minimumReservableCapacity>
  <vlanRangeAvailability>3747-3749</vlanRangeAvailability>
</link>
```

Delegated VLANs

- This work has been prototyped by Tom Lehman and Xi Yang of MAX
- This can be improved and expanded over time.

LINK#1				
InterfaceA			**InterfaceZ**	
Description	**ComponentId**	**AggregateManagerURL**	**Description**	**ComponentId**
BBN GPO-IG to Internet2 ION via NOX	urn:publicid:IDN+instageni.gpolab.bbn.com+interface+procurve2:5.24	https://www.instageni.gpolab.bbn.com:12369/protogeni/xmlrpc/am/2.0	Internet2 ION to BBN GPO-IG via NOX	urn:publicid:IDN+ion.internet2.edu+interface+rtr.newy:ae0:bbn
Intermediate Networks				
urn:publicid:IDN+nox.org				
Capacity	**MaximumReservableCapacity**	**MinimumReservableCapacity**	**Granularity**	**AvailableVlanRange**
1000Mbps	1000Mbps	1Mbps	n/a	3747-3749

Flack Stitching Support

- Rough support there now
- More to automate
- Add error handling
- Current procedure
  - Include ION in Aggs where you get the Ad
  - Create a link
  - Click ‘I’ and make it a stitched link
  - Click ‘Get Stitching Info’ (calls SCS)
  - Make reservation

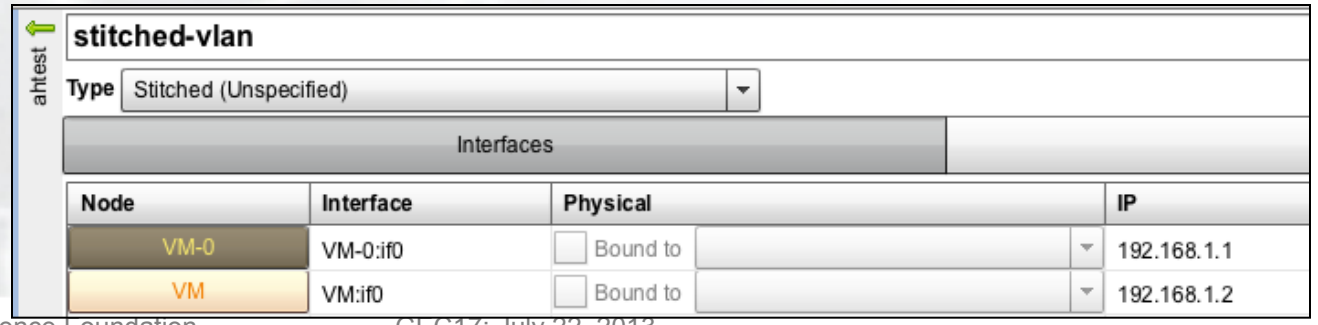

Aggregate Stitch Handling

- Read Request and Stitching Extension
  - Is (a) VLAN available?
  - Is node available?
- Handle Request
  - Configure the node
  - Configure the switch to connect the VLAN to the node
- Book-keeping
  - Node and VLAN are taken
- Return manifest
  - VLAN is reported in stitching extension

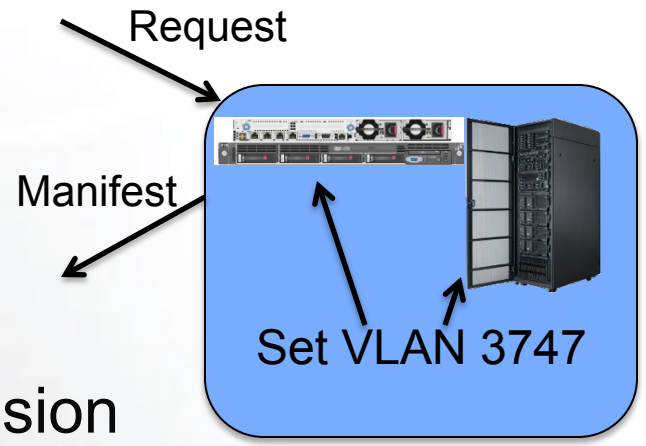

Stitching: What Next?

- Testing Continues
  - Detailed Test Plan
  - Early network and aggregate partners helping
- More Sites
  - After testing, all GENI racks
- More Providers
  - NLR, regionals, and campuses

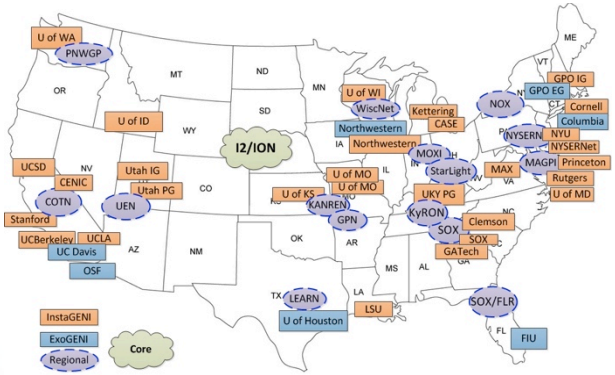

- More VLANs
  - VLAN ranges are small now – a scarce resource
  - GPO & rack teams will work to expand these
- Monitoring
  - VLAN use and circuit availability

Other Stitching for GEC18

- Stitcher improvements, new Omni 2.4
- Flack rollout
- Hardening SCS
- VLAN Delegation process improvements
- AL2S?

Stitching Enhancements

- Multi-point VLAN stitching
  - Stitching Extension Schema to represent multipoint L2/ VLAN bridging
  - SCS enhancements
    - Compute bridge paths
    - Support multi-point workflows
  - Work with EG, IG, I2 AL2s to add support
- Circuit Negotiation
  - Support negotiated tags with Aggregates
  - AM API v3 / 2-phase Allocate/Provision support

QUESTIONS? COMMENTS?

SLIDES FROM PLENARY

Tool's Job

- No central GENI stitching authority
  - Stitching is just a set of resource reservations at multiple aggregates.
- Any tool can do this.
  - The demo used an Omni script called 'stitcher.py'.
  1. Expand your request to find a path for your circuit
  2. Generate a request RSpec for each aggregate and make the reservations
  3. Check if any dynamic circuits were successfully created
  4. Report back a combined summary of what you have at all the aggregates

Handle Aggregate Response

- Tool reads VLAN out of manifest
- VLAN is inserted into request at next aggregate
  - GPO IG picks tag 3747, so request to ION uses 3747

```
<link client_id=”mylink">
  <component_manager…
  <component_manager…
…
<stitching>
  <path id=“mylink”>
    <hop id=“switch1:port1”>
      <vlan>3747</vlan>
```

- Tool handles many problems
  - VLAN is in use? Tool should try another
  - No VLANs available here? Try a different path
  - Something else, like no node available? Tell the user
- Stitcher checks for and handles these things.

3. Check Dynamic Circuits

- Dynamic Circuits take Time
  - OSCARS is configuring all your routers
  - This is fancy stuff
- Circuits can fail
- Tool can check and retry if it fails
  - Must wait to let OSCARS keep up

- Thankfully failures aren't that common

Stitching Extension

- Stitching RSpec Extension describes connections, paths, and requested or allocated VLANs

- This is how tools and aggregates talk about stitched circuits

  http://www.geni.net/resources/rspec/ext/stitch/

```
<stitching>
  <path id="mylink">
    <hop id="1">
      <link id=switch1:port1">
        ….
        <vlan…>3747</vlan…>
    <hop id="2">….
```

Stitching Extension

- In Aggregate advertisements
  - Local switch ports and VLAN ranges, and the remote switch port they connect to
- The SCS adds to your request a stitching extension
  - For each circuit, a series of switch ports / VLAN tags (hops) requested
- Manifests:
  - Your request with the VLAN tags assigned for each hop

```
<stitching>
 <path id="mylink">
  <hop id="1">
   <link id=switch1:port1">
    ….
    <vlan…>3747</vlan…>
  <hop id="2">….
```