GENI Stitching: Networks On-Demand

GEC17

July 22, 2013

Heidi Picher Dempsey

Luisa Nevers

Sponsored by the National Science Foundation

GENI and Stitching

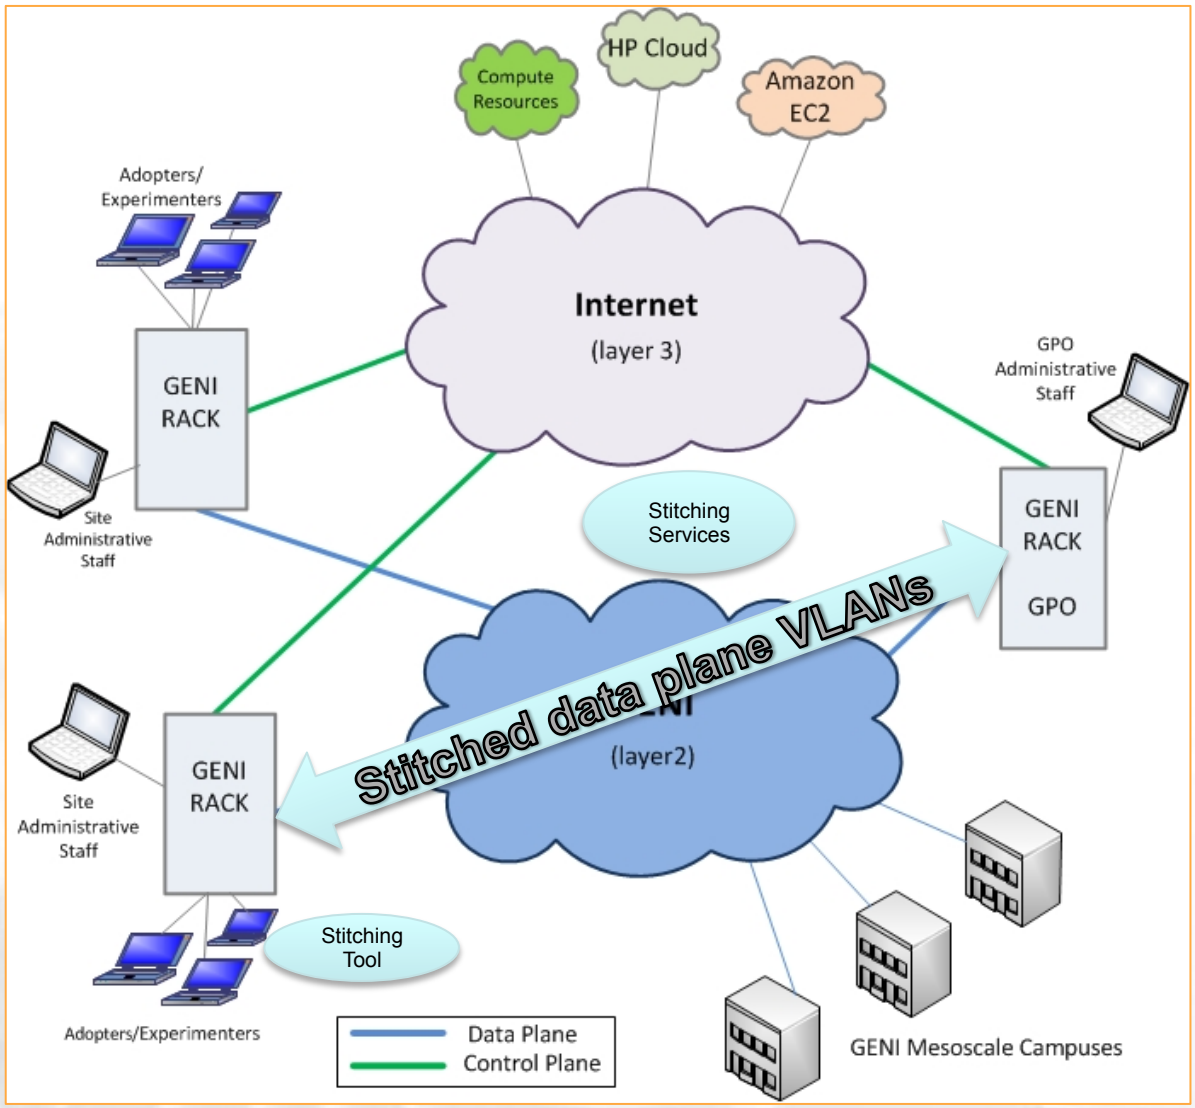

- GENI infrastructure is built on shared resources
- Providers only own part of an end-to-end data path
- GENI participants provide resource managers and tools to ease collaboration
- *Easier* experiment setup
- *Less* custom engineering

Stitching Background

- Stitching long supported in many commercial and R&E networks (e.g. OSCARS, GLIF)
- GENI rack teams both supported native stitching between like racks since earliest deployments and still do (ProtoGENI Flack and ORCA Flukes)
- GENI mesoscale supported custom-engineered VLAN connections for specific experiments
- GENI adopted interoperable stitching framework 2011

Stitching Today and Tomorrow

- Today we can stitch between unlike components (IG and EG racks, core, regional networks)

- Operations trials underway with I2 ION stitching aggregate, early racks, and brave regionals.
- Expect to expand to wherever GENI racks are deployed (note: international will require federation)
- Can easily go beyond GENI (open source, standards-based)

Stitching Collaborators Today

- GENI rack teams (RENCI/IBM, University of Utah/HP/ Northwestern)
  - Support GENI AM API and RSpec for stitching
  - Develop code to configure rack switches and assign slices

Stitching Collaborators (cont.)

Chad Kotil

Xi Yang

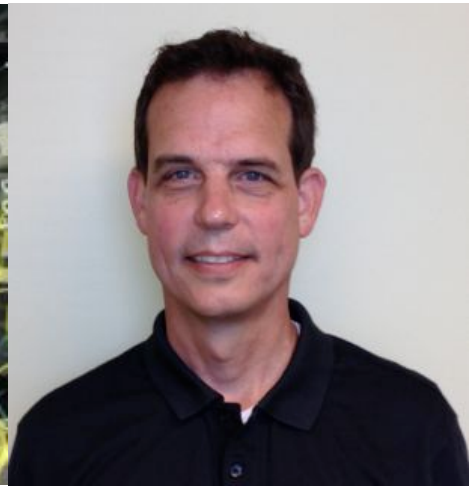

Tom Lehman

- Stitching Computation Service and ION Aggregate Teams (MAX/University of Maryland and Internet2)
  - Compute paths and dependencies
  - Request/Translate connections
  - Set up dynamic layer 2 data plane connections

Stitching Collaborators (cont.)

- Resource Owner/Operators (COTN/CENIC, KyRON, NOX, Stanford, UEN, University of Kentucky, University of Maryland, University of Utah)
  - Delegate VLANs and manage infrastructure

Northern Crossroads

Advanced networking for research and education in New England

The Corporation for Education Network Initiatives in California • 16700 Valley View Ave. Ste 400 • La Mirada, CA 90638

Stitching Collaborators - GPO

- Stitching client (using OMNI in the demo, other tools work too)
  - Request connections
  - Report results to experimenter
- GENI Clearinghouse (credentials only)
- Operations (GPO for initial trials)
  - Monitoring, reporting tools coming
  - Service desk help after that

Initial Operations Trials

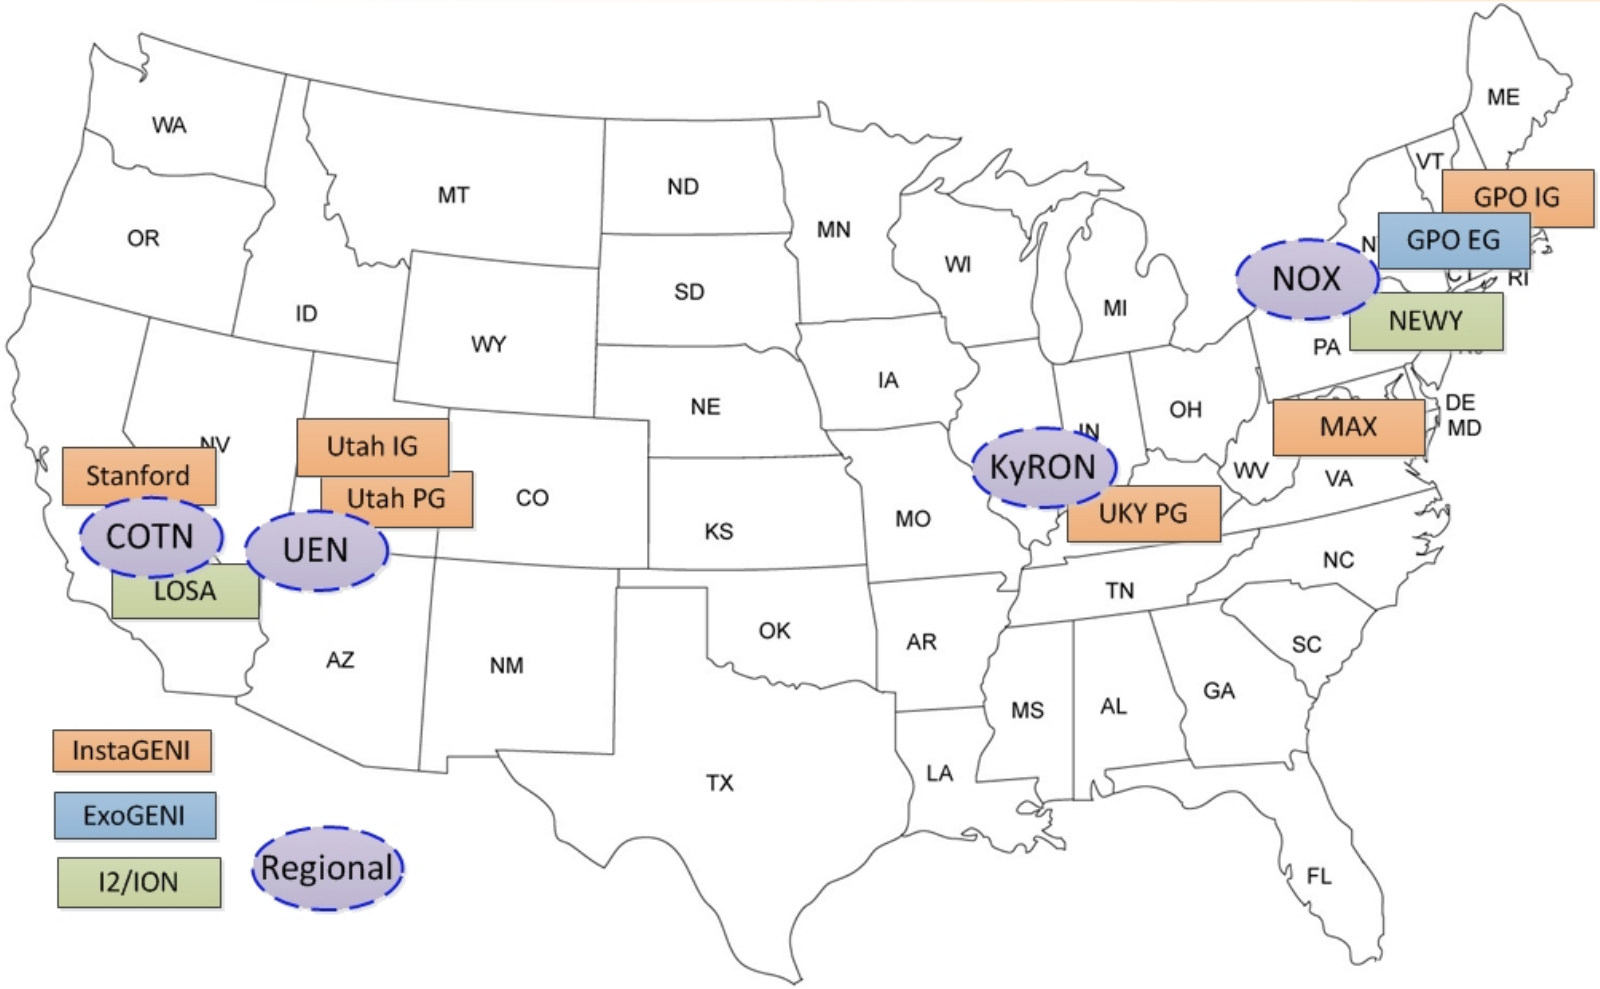

During operations trials, network engineers and brave experimenters exercise stitching “in the wild.” Developers collaborate and improve the tools and procedures.

Planned by Next GEC

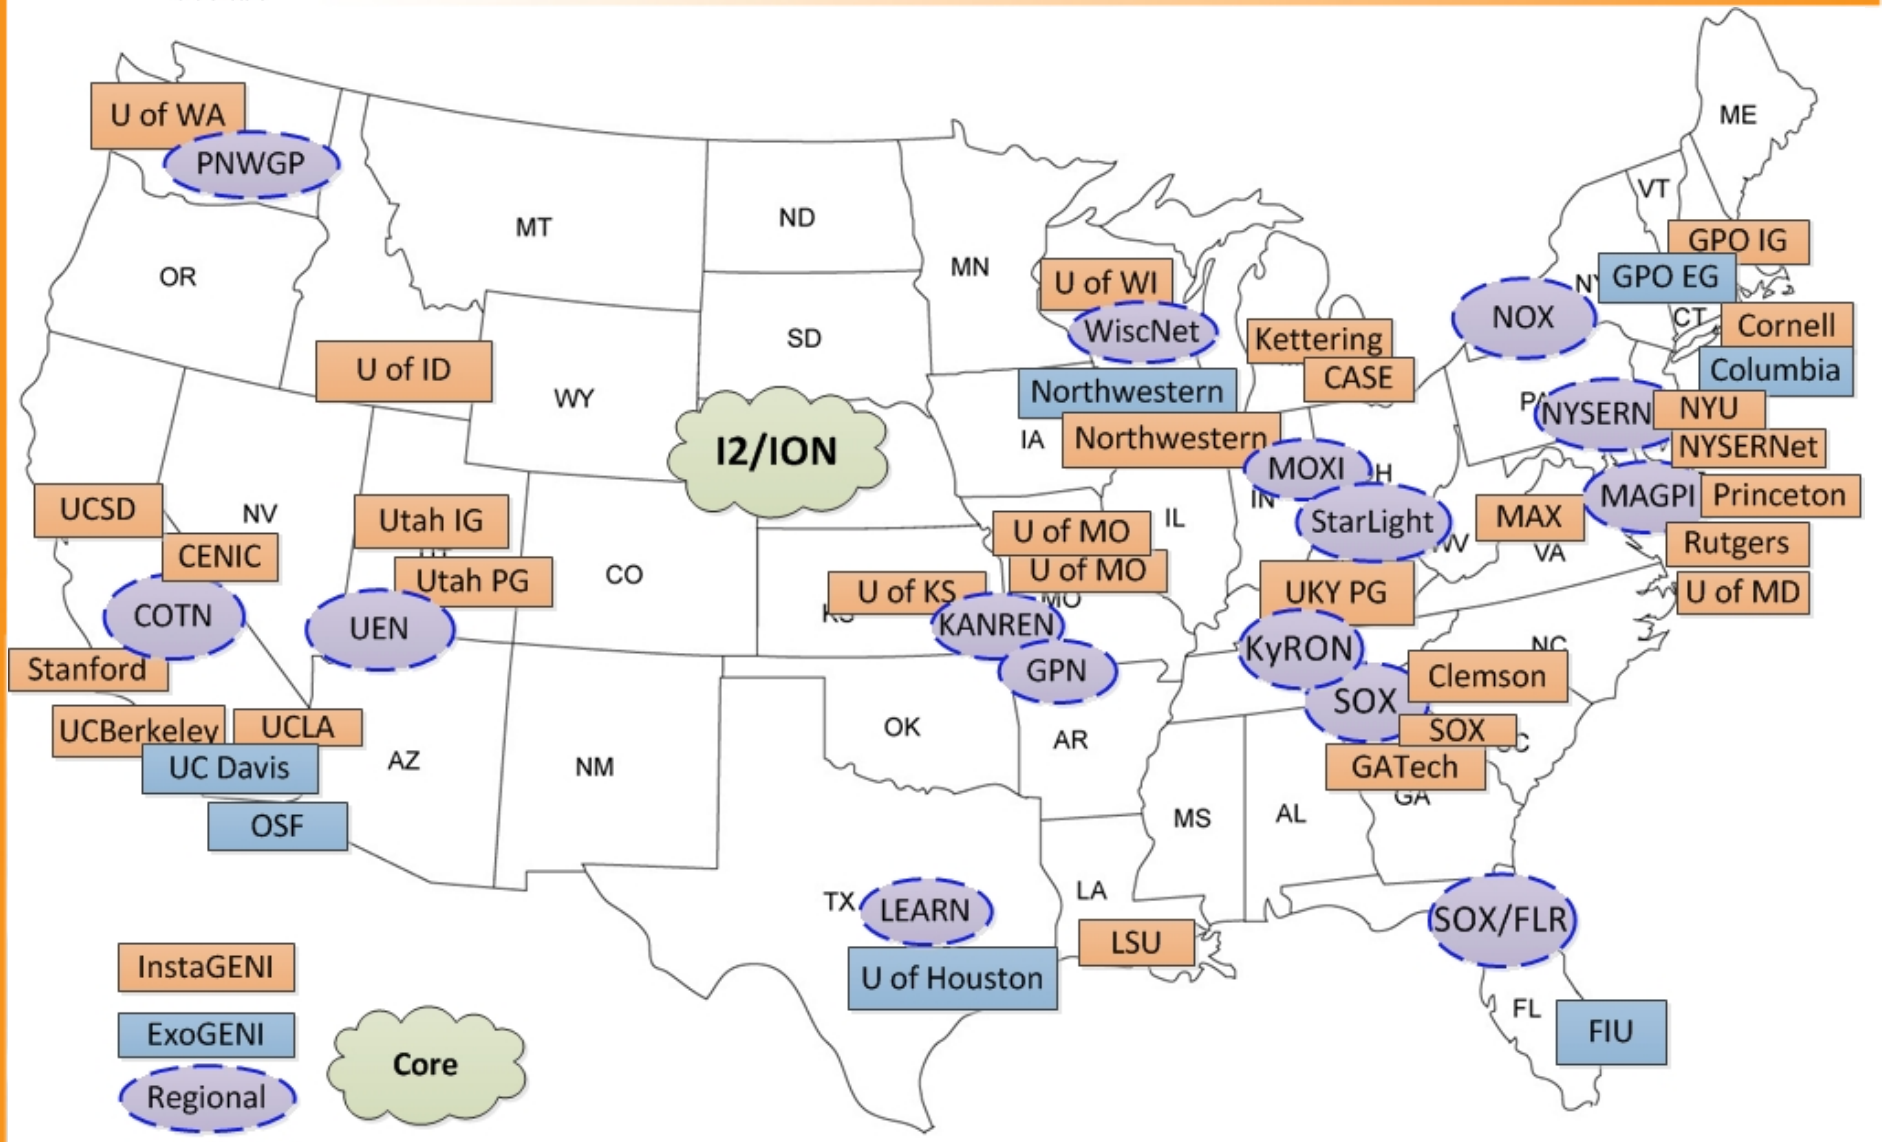

GENI Network Stitching Testing

- Functional, operator, resource, measurement and topology tests
- Exercise stitching tools in real operational environments
- Real experimenters and operators
- Find scaling issues and needed improvements
- Prepare for larger GENI rollout

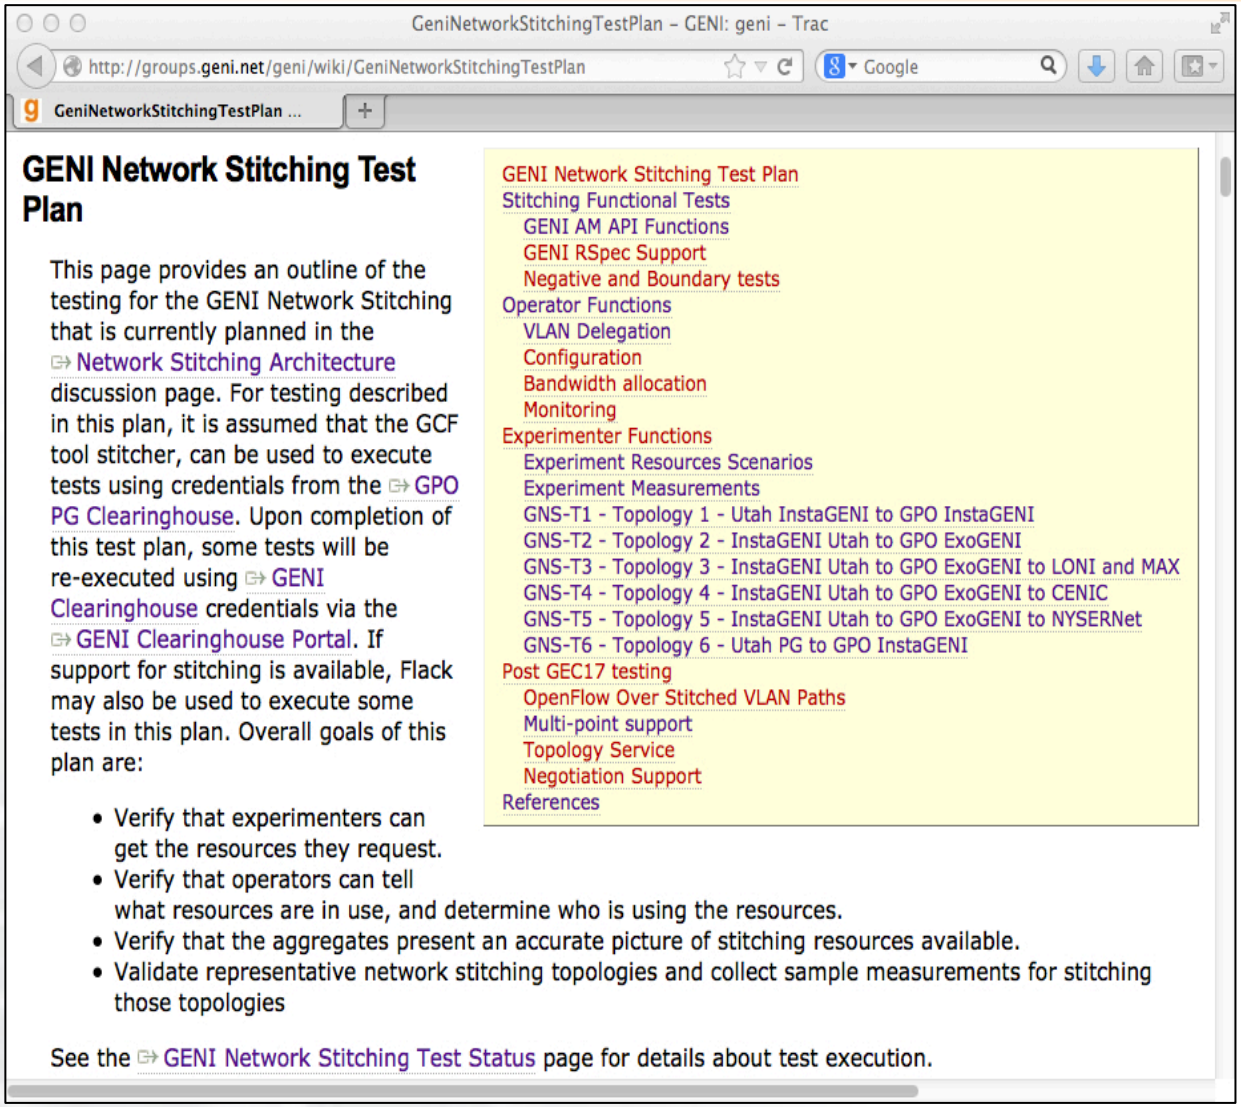

http://groups.geni.net/geni/wiki/GeniNetworkStitchingTestPlan

Stitching Testing Status

- Overall test status on GENI wiki (link below).
- Detailed test case captures linked to each test type (includes RSpecs).

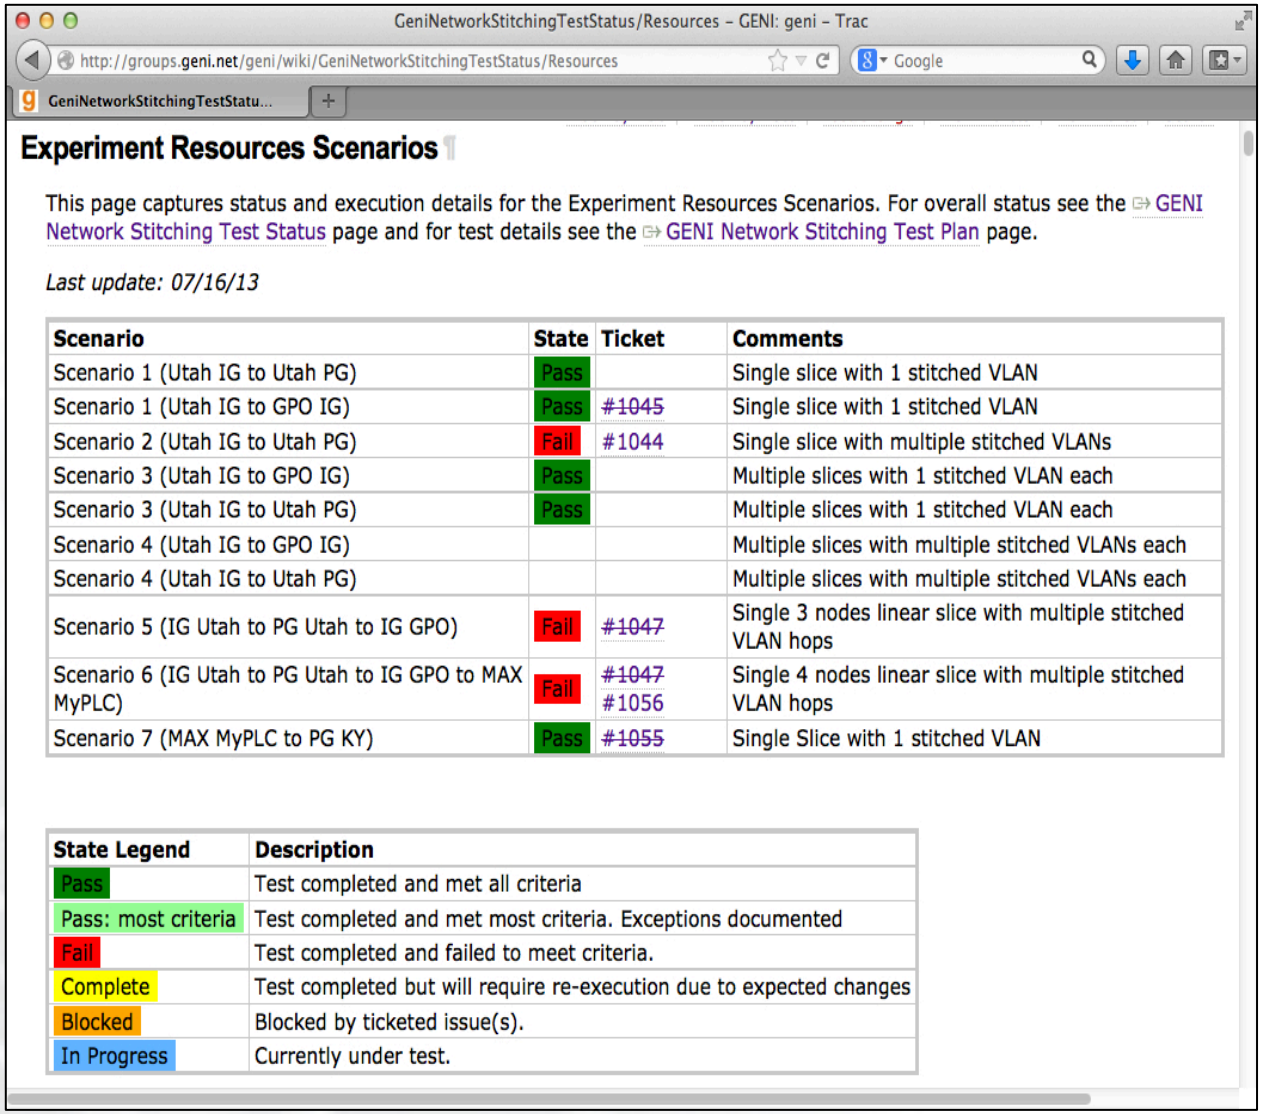

GeniNetworkStitchingTestStatus/Resources – GENI: geni – Trac

http://groups.geni.net/geni/wiki/GeniNetworkStitchingTestStatus/Resources

GeniNetworkStitchingTestStatu...

Experiment Resources Scenarios

This page captures status and execution details for the Experiment Resources Scenarios. For overall status see the GENI Network Stitching Test Status page and for test details see the GENI Network Stitching Test Plan page.

Last update: 07/16/13

Scenario	State	Ticket	Comments
Scenario 1 (Utah IG to Utah PG)	Pass		Single slice with 1 stitched VLAN
Scenario 1 (Utah IG to GPO IG)	Pass	~~#1045~~	Single slice with 1 stitched VLAN
Scenario 2 (Utah IG to Utah PG)	Fail	#1044	Single slice with multiple stitched VLANs
Scenario 3 (Utah IG to GPO IG)	Pass		Multiple slices with 1 stitched VLAN each
Scenario 3 (Utah IG to Utah PG)	Pass		Multiple slices with 1 stitched VLAN each
Scenario 4 (Utah IG to GPO IG)			Multiple slices with multiple stitched VLANs each
Scenario 4 (Utah IG to Utah PG)			Multiple slices with multiple stitched VLANs each
Scenario 5 (IG Utah to PG Utah to IG GPO)	Fail	~~#1047~~	Single 3 nodes linear slice with multiple stitched VLAN hops
Scenario 6 (IG Utah to PG Utah to IG GPO to MAX MyPLC)	Fail	~~#1047~~ #1056	Single 4 nodes linear slice with multiple stitched VLAN hops
Scenario 7 (MAX MyPLC to PG KY)	Pass	~~#1055~~	Single Slice with 1 stitched VLAN

State Legend	Description
Pass	Test completed and met all criteria
Pass: most criteria	Test completed and met most criteria. Exceptions documented
Fail	Test completed and failed to meet criteria.
Complete	Test completed but will require re-execution due to expected changes
Blocked	Blocked by ticketed issue(s).
In Progress	Currently under test.

http://groups.geni.net/geni/wiki/GeniNetworkStitchingTestStatus

Stitching Variations

GENI Network Stitching combinations supported and tested:

	GPO InstaGENI	Utah InstaGENI	Utah PG	UKY PG	MAX
GPO InstaGENI	N/A	✔	✔	✔	✔
GPO ExoGENI	Not Supported	✔	Not tested	Not tested	Not tested
Utah InstaGENI	✔	N/A	✔	✔	✔
Utah PG	✔	✔	N/A	✔	✔
UKY PG	✔	✔	✔	N/A	✔
MAX	✔	✔	✔	✔	✔
Stanford Non-GENI	✔	Not tested	Not tested	Not tested	Not tested

How to Stitch: Operators

- Select infrastructure for stitching.
  - VLANs, LANs, racks, other aggregates
  - ION core stitching only for GENI trials
- Select local endpoints and available bandwidths
- Delegate VLANs for GENI stitching.
- Set up additional local VLANs (optional).
- Run a GENI stitching service (optional).
- Contact help@geni.net to integrate, test, and announce.

Stitching VLAN Delegation

- Internet 2 ION can allocate any VLAN ID dynamically in core network and provide VLAN translation.
- Most network providers and campuses delegate smaller VLAN ID ranges for GENI.
- VLAN Stitching Computation Service combines ION and delegated VLAN segments to find paths endpoints.
- Operations trials public wiki page records who delegated what:

https://hpn.dragon.maxgigapop.net/twiki/bin/view/GENI/StaticNetworksView

Delegated VLANs Page

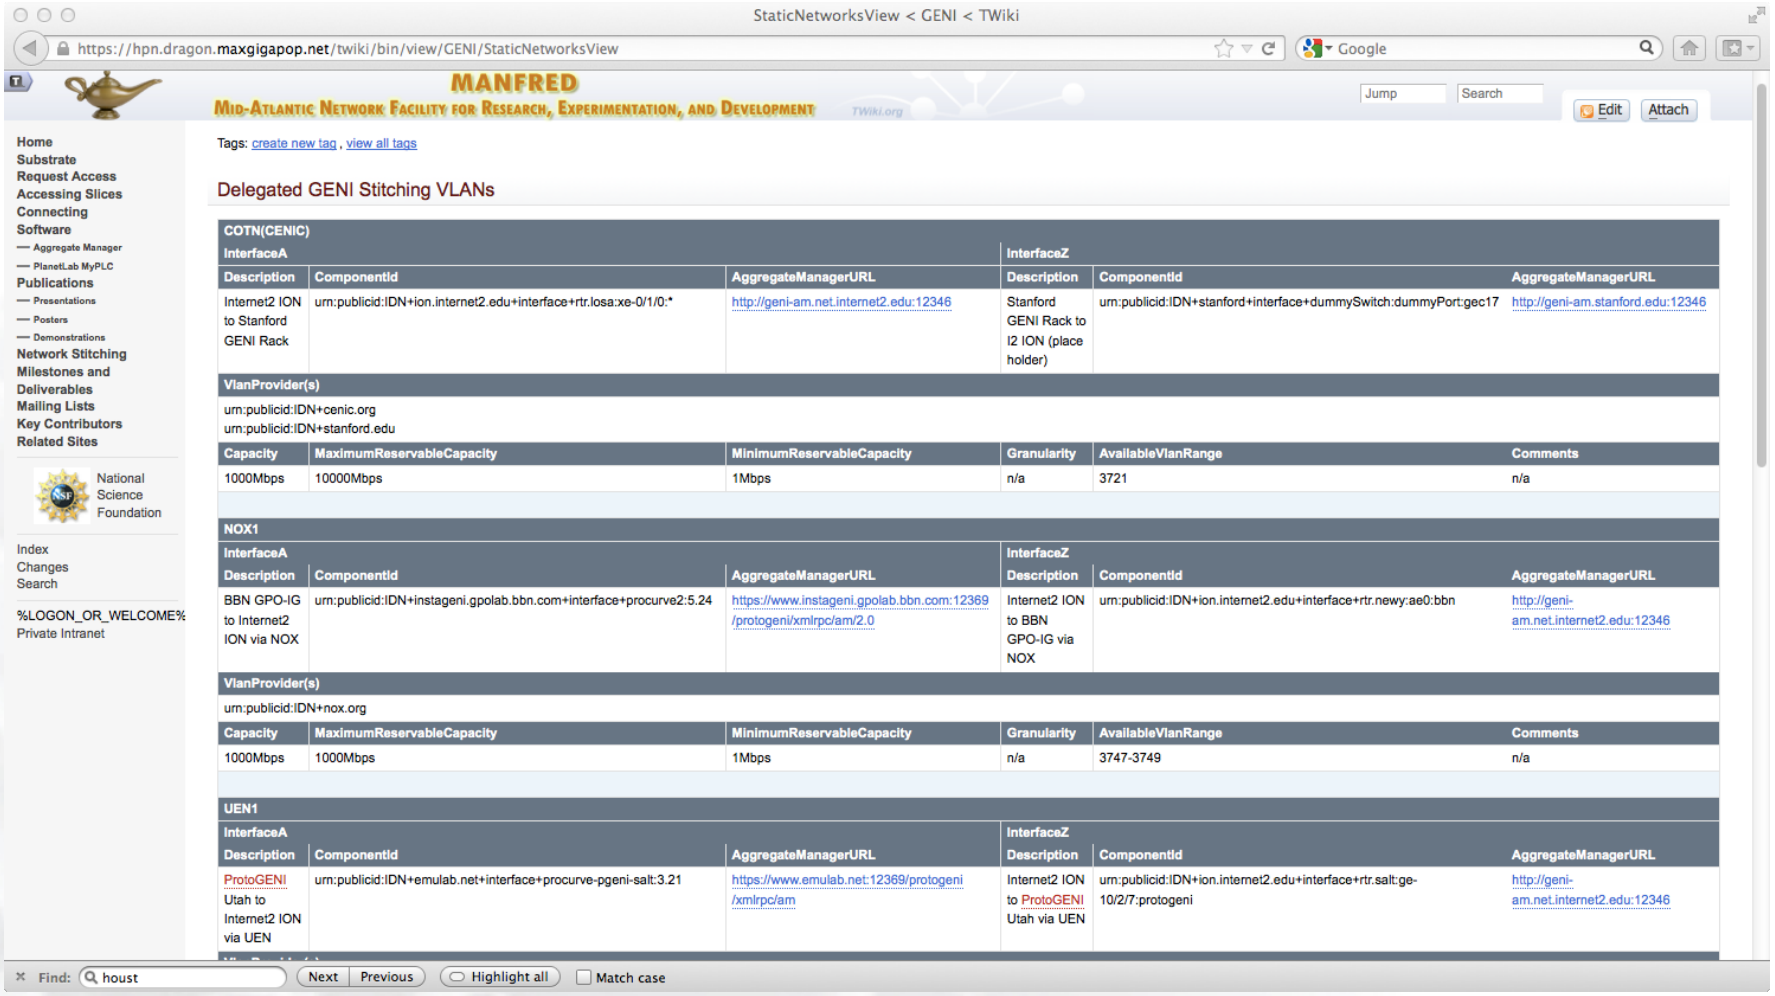

How to Stich: Experimenters

1. Get Omni Client

- GENI network stitching support introduced in gcf 2.3.1
- New version under test -- check GCF wiki page for availability and features:

  http://trac.gpolab.bbn.com/gcf/wiki/WikiStart
- GENI network stitching support also in progress for Flack
- GENI stitching via GENI Portal in development, not yet available

How to Stich: Experimenters

2. Find sites that support GENI Stitching

- Test sites available now, more coming after GEC17.
- Current sites that have successfully passed stitching tests:
  - GPO (InstaGENI and ExoGENI)
  - University of Utah (InstaGENI and ProtoGENI)
  - University of Kentucky (ProtoGENI)
  - MAX (MyPLC)
  - CENIC (with Stanford non-GENI Compute Resource)

How to Stich: Experimenters

3. Get request RSpecs for your endpoints

- GENI Network Stitching only requires a *component manager id* in the <*link*> definition for each of the end-point aggregates:

```
<link client_id="link”>
<component_manager name="urn:publicid:IDN+instageni.gpolab.bbn.com+authority+cm”/>
<component_manager name="urn:publicid:IDN+emulab.net+authority+cm"/>
….
</link>
```

- Example Rspecs in experiment on GENI wiki:

http://groups.geni.net/geni/wiki/GENIExperimenter/ExperimentExample-stitching

How to Stich: Experimenters

4. Execute Experiment

- Simple experiment set-up commands. Code in *stitcher.py* creates slivers at all needed aggregates, including network stitching path.
- *Stitcher.py* parses your RSpec to determine ***all*** compute resource aggregates. No longer need to specify "*-a aggregateName*".
- *Stitcher.py* uses the MAX Stitching Computation Service (SCS) to determine VLAN path needed for your experiment

How to Stich: Experimenters

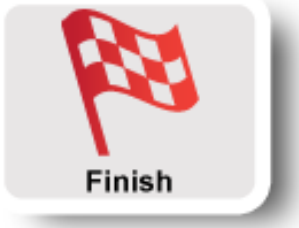

5. Finish Experiment

- Simple experiment tear-down commands.
- S*titcher.py* deletes your slivers at all known slice aggregates, including your network stitching path.
- *Stitcher.py* deletes slivers at all known aggregates on the GENI Portal, if it does not know which aggregates are part of your slice.

Demo GENI Network Stitching

The GENI Network Stitching live demo shows an experiment that creates a network stitching path between the GPO InstaGENI and the Utah InstaGENI. Shown below is the end-to-end network path:

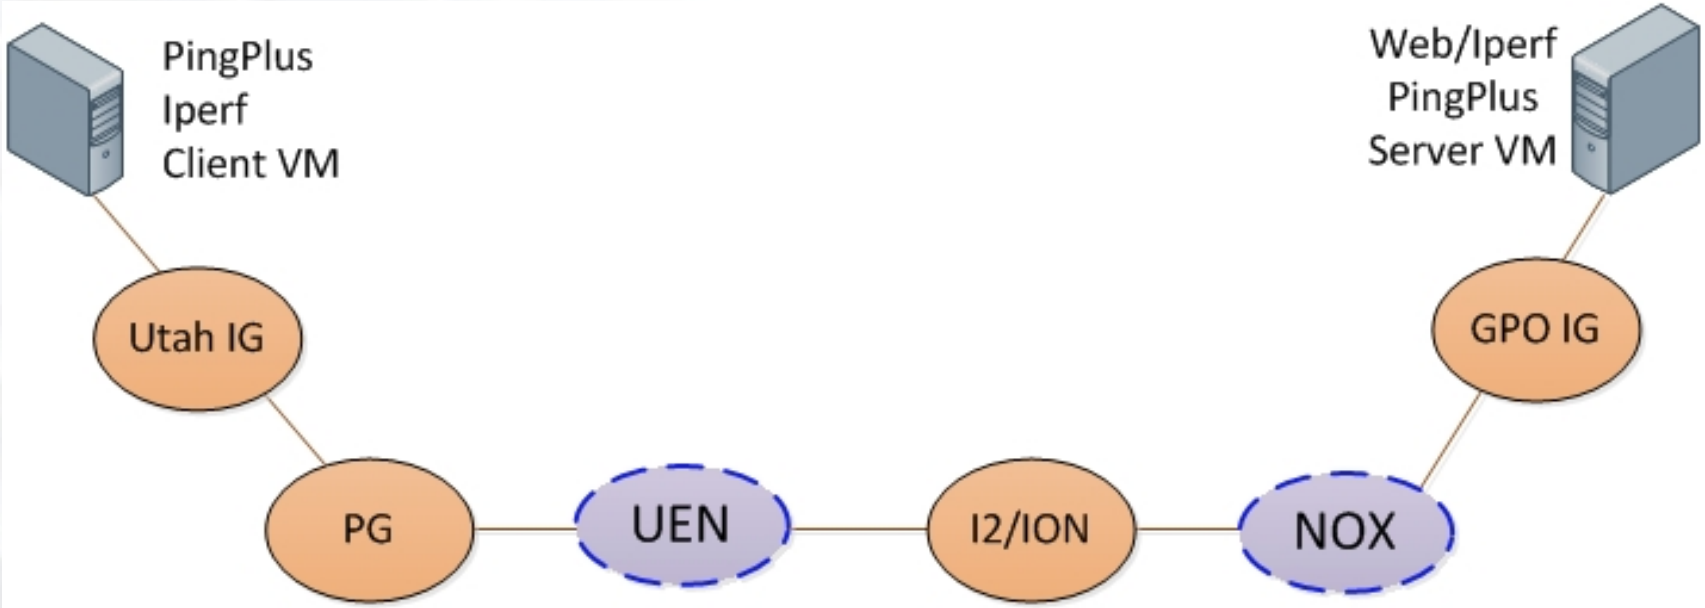

Note: UEN and NOX have delegated VLANs in the path

Iperf, wget and PingPlus traffic is exchanged in the experiment.

Demo GENI Network Stitching (cont.)

Get Request RSpec

- RSpec used for this demo is available at:

  http://groups.geni.net/geni/browser/trunk/stitch-examples

- Remember the portion of the RSpec that matters to stitching:

<link client_id="link">

<component_manager name="urn:publicid:IDN+instageni.gpolab.bbn.com+authority+cm"/>

<component_manager name="urn:publicid:IDN+utah.geniracks.net+authority+cm"/>

….

</link>

Demo GENI Network Stitching (cont.)

Stitcher Workflow:

1. **Stitcher gets your slice credential** and verifies that it is a valid slice in which to create your circuit and runs rspeclint on your RSpec file, if rspeclint is found in search path.
2. **Stitcher.py contacts SCS** to get the list of AMs needed to establish the stitching path, including, if needed, the Internet2 ION aggregate.
3. **Stitcher generates request RSpec** for each of the aggregates in the network path requesting both network and compute resources. Slivers are created at each aggregate.
4. **Stitcher.py requests ION DCN circuits** and checks for circuit readiness. It may try up to 10 times, pausing 30 seconds between each try. Usually needs only one try.
5. **Stitcher.py assembles a combined manifest RSpec** which shows all aggregates and VLANs used in the slice.

Demo GENI Network Stitching (cont.)

- What does it look like?

$ omni.py createslice *sliceName*

$ stitcher.py createsliver *sliceName rspec.xml* –o

$ readyToLogin.py *sliceName* –a *AM*

…. the experiment

$ stitcher.py deletesliver *sliceName*

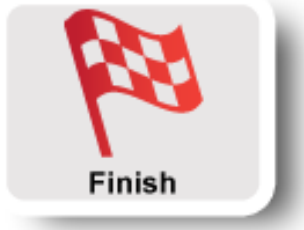

Stitching Variations

Multiple stitched connections share the infrastructure. Running alongside this live demo are several other live slices using stitched topologies:

- GPO ExoGENI to Utah InstaGENI
  - A previously set up slice named *Indemo1* is running between GPO EG and Utah IG using a network stitching path.
  - Traffic exchange between GPO ExoGENI (192.1.242.20) to Utah InstaGENI (pc1.utah.geniracks.net port 32058) via stitched path.
  - Let's see the live results
- GPO InstaGENI to Stanford via CENIC's COTN
  - A previously set up slice named *Instan* is running between GPO InstaGENI and a non-GENI Stanford compute resource using a stitching path via COTN.
  - Traffic exchange between GPO VM (pc5.utah.geniracks.net port 31290) to Stanford host (netw-cenic01.stanford.edu)

Stitching Varied Compute Nodes

- Max MyPLC to University of Kentucky PG
  - A previously set up slice named *max-pg-ky* is running between MAX MyPLC and University of Kentucky PG using a network stitching path.
  - Traffic exchange between MAX MyPLC host (planetlab2.dragon.maxgigapop.net) to UKY PG (pc63.uky.emulab.net port 33338) via stitched path.

Stitching Test Status

Testing started May 22, 2013:

- Survey of AM API support, RSpecs support, Negative and Boundary testing completed.
- Network topology tests ran successfully for current aggregates. More interesting scenarios coming as we add more aggregates and VLANs.
- Reliability has progressed noticeably since testing started.
- Some known issues remain, but workarounds are available.
- Features not yet available: OpenFlow over stitched VLAN paths and multipoint VLAN topologies

Stitching Test Findings

Issues found for all aggregates; most addressed.

- ION/I2 issues
  - ION DCN Circuits allocation/creation
  - SCS configuration
  - Advertisement/Manifest RSpecs content
  - Some scaling/speed issues
- InstaGENI/PG issues
  - Utah InstaGENI endpoint requires capacity be for link
  - Duplicate IP addresses with multiple links
  - Requesting invalid capacity does not generate error
  - Stitching slivers /etc/hosts incorrect remote entries
  - Default MTU failures with VLAN tags

Stitching Going Forward

- Operations trials continuing
- If you are getting a GENI rack, you'll be in the trials automatically
- Open to all interested providers, experimenters, and developers
- Join us! help@geni.net

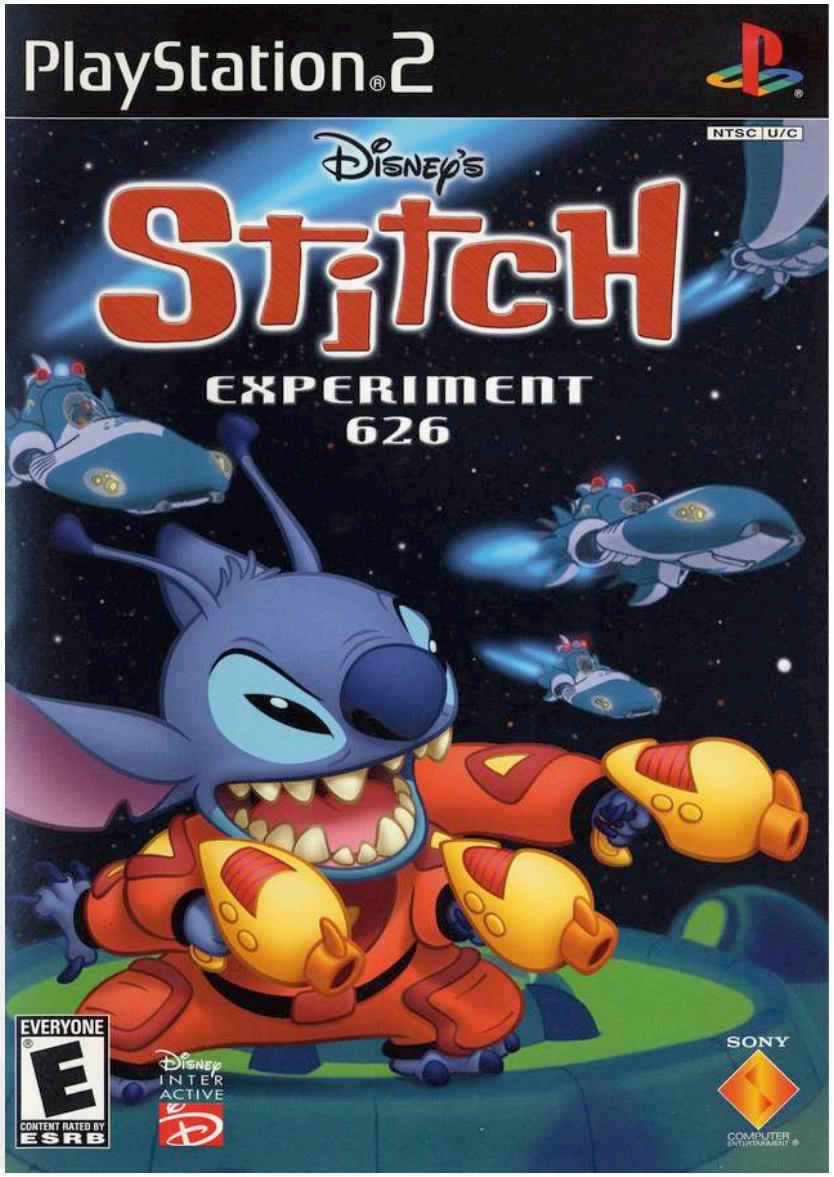

Apologies to Disney and Sony

Backup

Stitching Between Different Rack Types

GPO ExoGENI to Utah InstaGENI Stitching

lnevers — lnevers@ig-utah:~ — ssh — 109×19

```
lnevers@sendaria:
lnevers@sendaria:~$ ssh -p 33082  -i /home/lnevers/.ssh/geni_cert_portal_key lnevers@pc5.utah.geniracks.net
Last login: Tue Jul 16 11:20:20 2013 from sendaria.gpolab.bbn.com
[lnevers@ig-utah ~]$ ping -c 5 192.168.1.2
PING 192.168.1.2 (192.168.1.2) 56(84) bytes of data.
64 bytes from 192.168.1.2: icmp_req=1 ttl=64 time=0.024 ms
64 bytes from 192.168.1.2: icmp_req=2 ttl=64 time=0.020 ms
64 bytes from 192.168.1.2: icmp_req=3 ttl=64 time=0.020 ms
64 bytes from 192.168.1.2: icmp_req=4 ttl=64 time=0.021 ms
64 bytes from 192.168.1.2: icmp_req=5 ttl=64 time=0.014 ms

--- 192.168.1.2 p
5 packets transmi
rtt min/avg/max/m
[lnevers@ig-utah
```

```
lnevers@sendaria:~$ ssh -i /home/lnevers/.ssh/geni_cert_portal_key root@192.1.242.10
Linux debian 2.6.32-5-amd64 #1 SMP Mon Jun 10 10:22:28 UTC 2013 x86_64

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Wed Jul 17 17:41:05 2013 from sendaria.gpolab.bbn.com
root@eg-gpo:~# ping -c 5 192.168.1.1
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp_req=1 ttl=64 time=0.018 ms
64 bytes from 192.168.1.1: icmp_req=2 ttl=64 time=0.032 ms
64 bytes from 192.168.1.1: icmp_req=3 ttl=64 time=0.024 ms
64 bytes from 192.168.1.1: icmp_req=4 ttl=64 time=0.027 ms
64 bytes from 192.168.1.1: icmp_req=5 ttl=64 time=0.021 ms

--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3999ms
rtt min/avg/max/mdev = 0.018/0.024/0.032/0.006 ms
```

Stitching to Endpoints Outside GENI

GPO InstaGENI to Stanford via CENIC's COTN Stitching

```
lnevers@sendaria:~$ ssh -p 36154  -i /home/lnevers/.ssh/geni_cert_portal_key lnevers@pc2.instageni.gpolab.bbn.com
Last login: Tue Jul 16 13:32:20 2013 from sendaria.gpolab.bbn.com
[lnevers@ig-gpo1 ~]$ ping -c 5 192.168.1.1
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp_req=1 ttl=64 time=80.3 ms
64 bytes from 192.168.1.1: icmp_req=2 ttl=64 time=80.3 ms
64 bytes from 192.168.1.1: icmp_req=3 ttl=64 time=80.3 ms
64 bytes from 192.168.1.1: icmp_req=4 ttl=64 time=80.3 ms
64 bytes from 192.168.1.1: icmp_req=5 ttl=64 time=80.3 ms

--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4005ms
rtt min/avg/max/mdev = 80.322/80.367/80.399/0.255 ms
[lnevers@ig-gpo1 ~]$
```

```
lnevers@sendaria:~$ ssh netw-cenic01.stanford.edu -i /home/lnevers/.ssh/id_rsa
Last login: Tue Jul 16 10:18:34 2013 from 128.89.91.19
[lnevers@netw-cenic01 ~]$ ping 192.168.1.2 -c 5
PING 192.168.1.2 (192.168.1.2) 56(84) bytes of data.
64 bytes from 192.168.1.2: icmp_seq=1 ttl=64 time=80.3 ms
64 bytes from 192.168.1.2: icmp_seq=2 ttl=64 time=80.4 ms
64 bytes from 192.168.1.2: icmp_seq=3 ttl=64 time=80.3 ms
64 bytes from 192.168.1.2: icmp_seq=4 ttl=64 time=80.3 ms
64 bytes from 192.168.1.2: icmp_seq=5 ttl=64 time=80.3 ms

--- 192.168.1.2 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4084ms
rtt min/avg/max/mdev = 80.326/80.358/80.418/0.360 ms
[lnevers@netw-cenic01 ~]$
```

Stitching PlanetLab and Emulab Nodes

MAX MyPLC to UKY PG Stitching