Developer Topics: Stitching

Stitching is Here; How can we improve it?

Aaron Helsinger, GPO

GEC19: March 19, 2014

Stitching Status

- 12 Aggregates support GENI VLAN Network Stitching
- > 5000 stitch requests since January 1
- Stitcher 2.5 just released
- People are using it, using up the VLANs, and pushing the boundaries

Stitching Issues

- But there are things we can improve…
- We're running out of resources
  - We need better usage information to ask for more VLANs and bandwidth
  - Monitoring effort on usage will help everyone
- We're stressing ION and regionals
  - What can we do? Will AL2S help?

What can we do?

Stitching Issues (2)

- Stitching takes too long
  - We need VLAN negotiation, accurate Ad RSpec availability information
- It's hard to tell why stitching fails
  - Need better error messages from AMs, tools, SCS
- Tool support is limited
  - Would a workflow service help? A service to find connected AMs?

What can we do?

Address Stitching Issues

- Stitching is slow
  - Ad RSpecs should reflect provisioned bandwidth and VLANs
    - So tools can guide requests that might succeed, start with VLAN tags that are available
  - VLAN tag negotiation
    - Allow tools to cheaply (& quickly) negotiate with AMs for an available VLAN tag
    - Requires Allocate/Provision from AM API v3
  - AL2S is faster – when will we be ready?
- Experimenters can't understand failures
  - Consistent error messages from aggregates would help

Address Issues (2)

- Running out of resources
  - Collecting information on usage, errors would help motivate requests
  - Monitoring effort will provide much of that data
  - We can then ask for more VLANs or bandwidth
- Limited tool support for stitching
  - Building support services would encourage tools
  - Build a service to find interconnected AMs
    - Tools can guide experimenters to make good requests
  - Workflow service to orchestrate stitching
    - The logic is complex for tool developers; outsource it
  - Add friendly stitching interfaces in GENI Desktop, Jacks, etc

Address Issues (3)

- ION errors and switch reconfiguration delays frustrate users
  - ION circuits fail after allocation with cryptic error messages
  - Allocating a circuit on ION requires rewriting multiple switch configurations – this takes time, causes other errors
  - The result is that circuits across ION take time
  - AL2S should make many of these things better

Missing Features

- Stitch to all nodes at an AM
  - Connect to all nodes at 2 connected AMs with 1 VLAN
    - Not just single node to single node
  - Can AMs do this?

- OpenFlow stitching
  - A VLAN whose routing is OpenFlow controlled
  - It is in the v2 stitching schema
  - Other options are possible, raised yesterday at the Programmable WAN session
  - Need AM and tool support

Missing Features (2)

- Common Stitch Point support
  - Support stitching to arbitrary other switch/ports
  - Is this a `node` in the Ad RSpec?
  - Or an `interface_ref` with a matching `external_ref`?

- Support Stitching Schema v2

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

  - Adds OpenFlow support, better capabilities, etc
  - Adopted ages ago, but never implemented
  - Can we agree on a cut-over to point to the new schema?

Missing Features (3)

- Multicast / Multi Point support
  - Stitching today gives us point-to-point VLANs
  - Can we get true multicast? Using AL2S?
  - Need SCS, aggregate, tool support
- AL2S AM
  - AL2S will be replacing ION eventually
  - How will we transition and support stitching?

Issues to Resolve

- Stitching Computation Service
  - We depend on it
  - The SCS needs replication, support
  - SCS calls are unauthenticated today
- Capacity units
  - Kbps in SCS, ION, InstaGENI, stitcher
  - Bps in ExoGENI
  - Can we standardize?
- Aggregates of Aggregates
  - The ExoSM, for example
  - How do tools know which AM to contact for a given resource? Can the experimenter control it?

Issues to Resolve (2)

- URN format conventions
  - /: MAX & ION use it to be consistent with OSCARS
    - It's odd in a URN, since public to URN translation escapes it
  - Structure of name of interfaces
    - Node:interface or node:interface:subinterface is typical
      - Human readable
      - Real world correlation

```
urn:publicid:IDN+ion.internet2.edu+interface
+rtr.kans:ge-10/2/9:protogeni
```

    - Alternative is free-form
      - Less config changes when devices change

```
urn:publicid:IDN+exogeni.net:bbnvmsite+interface
+Bbn:ExoGeni:TenGigabitEthernet:1:ethernet
```

    - Worth standardizing?

Issues to Resolve (3)

- Circuit expiration
  - If 1 segment of a circuit expires before another, AMs disagree on what VLANs are available – errors result
    - E.G. PG Utah expires the sliver, ION has still reserved the VLAN tag. Then PG Utah allocates the VLAN tag to someone else, who gets an error at ION
  - Must all VLAN slivers have a common initial lifetime?
  - Must tools renew all slivers to the same time?
  - Or do AMs need to negotiate an expiration time?

Discussion

- What would improve stitching the most?
- What are the biggest missing features?
- Where do we want to drive the effort?