Moving Towards SDN @ AmLight

4th Annual Global LambdaGrid Workshop

Oct 1st 2014

Queenstown, New Zealand

Jeronimo Bezerra <jbezerra@fiu.edu>

Who we are

AMPATH:

- GLIF GOLE and Academic IXP in Miami
- Interconnects all Latin America RENs to other NRENs in the world

SouthernLight

- GLIF GOLE and Academic IXP in Sao Paulo
- Interconnects all Brazilian RENs and RedCLARA to AmLight

AmLight:

- A set of trans-continental high-bandwidth links connecting SounthernLight, REUNA and AURA/Chile to AMPATH

Partners: Florida International University/FIU, NSF, ANSP, RNP, RedCLARA, REUNA and AURA

AmLight Today

A set of 4 x 10G links connecting South America to AMPATH

- Academic Layer 2 Ring (Miami-Sao Paulo-Santiago)
- MPLS Ring (Miami-Fortaleza-Rio-Sao Paulo)
- Later this year: 100G link between Sao Paulo and Miami (OpenWave)

Connections:

- 13 RENs
- > 1000 Universities and Research Centers

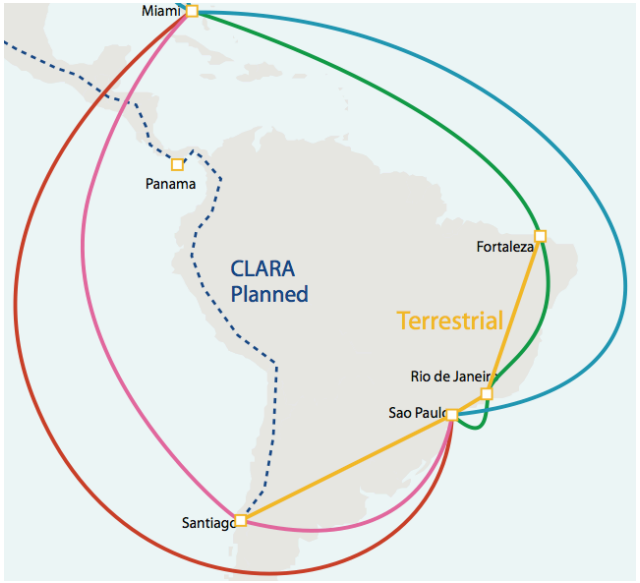

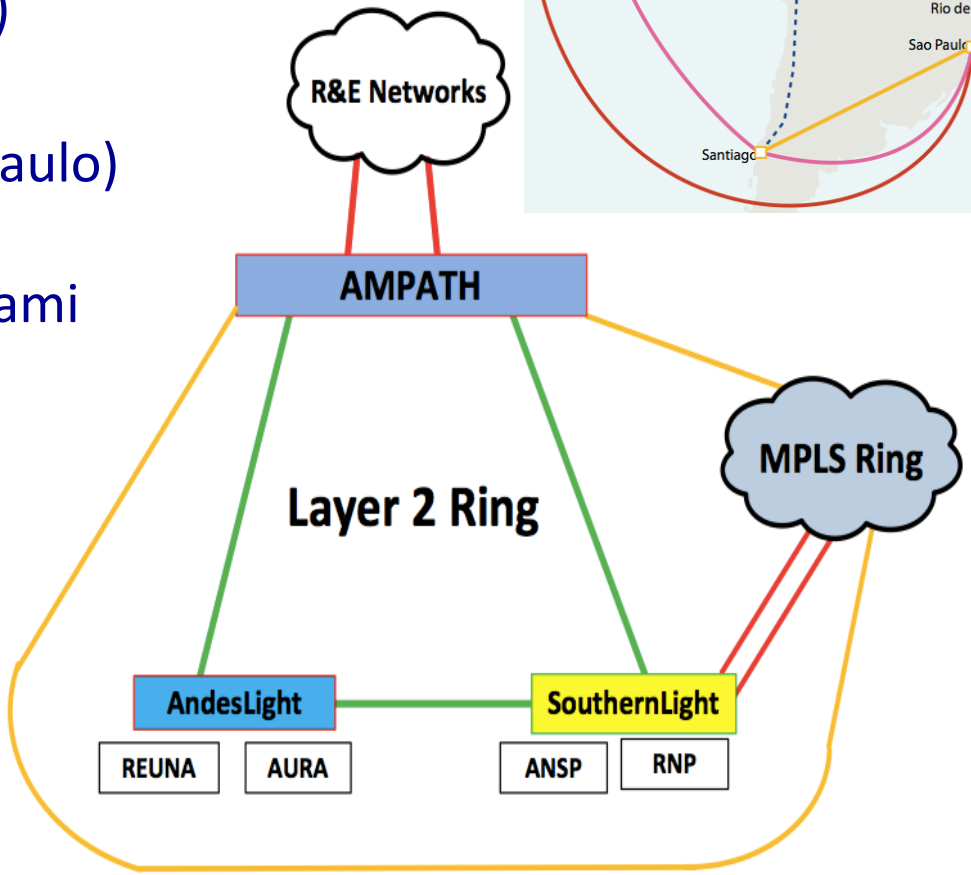

AmLight Layer 2 Ring
Up to two months ago

- Mostly used for academic/research collaboration
- Configuration based on static VLANs
- Multiple instances of Brocade per-VLAN RSTP
- Receives redundancy from Academic Layer 3 Ring:
  - IEEE 802.1ad (QinQ) deployed in Sao Paulo
  - 3x 10G ports dedicated to this redundancy
- Very resilient, 100% availability in 2013 (assuming minimum of 10Gbps)

If everything was working fine, why change it?

Why has AmLight moved towards SDN?

Key motivations:

Improving operations efficiency

Introducing network programmability

Motivation 01:
Improving Operations Efficiency

Amount of layer 2 circuits requested and networks involved makes the provisioning a complex process:

- Some circuits involve up to seven different networks
  - High level of coordination required with diverse network teams
- Multiple technologies involved
  - From Layer 1 to MPLS
- Some circuits took weeks or even months to be provisioned

It was necessary to find a way to improve this situation…

Motivation 02:
Introducing Network Programmability

- The lack of support for network programmability compromises network-aware demos and applications
- Researchers could only view the network status (SNMP)

How has AmLight addressed these motivations? (1)

- Is it possible to use SDN to address these issues?
- Is it the right moment for using SDN?
- Do my devices support Openflow?
- Are we going to lose features?

How has AmLight addressed these motivations? (2)

- Is it possible to use SDN to address these issues?
  - Yes!
- Is it the right moment for using SDN?
  - For AmLight, which only supports Layer 2, Yes!
- Do our devices support Openflow?
  - After diverse tests and simulations, we confirmed: Yes! 1.0
- Are we going to lose features?
  - Yes
    - Some modules don't support Openflow counters
    - LACP is not supported
  - But we have workarounds. Move forward!

Scenario Deployed (1)

- *Devices:*
  - *Brocade MLXe (8x10G)*
  - *Brocade XMR (4x10G)*
  - *Brocade CES*
    - *Version 5.6d*
- *Improving Operations Efficiency:*
  - *Internet2's OESS*
- *Adding Network Programmability:*
  - *Internet2's FlowSpace Firewall*
- *Final Approach:*
  - *Openflow 1.0*
  - *Sao Paulo has a "Hybrid" port*
  - *Chile is "full SDN"*
  - *Controllers installed in Miami, mirror in SP**
  - *Integration with Internet2 via OSCARS*
    - *NSI in the future*
  - *Non-Academic traffic remains "legacy/protected"*
  - *7 new 10G connections required*

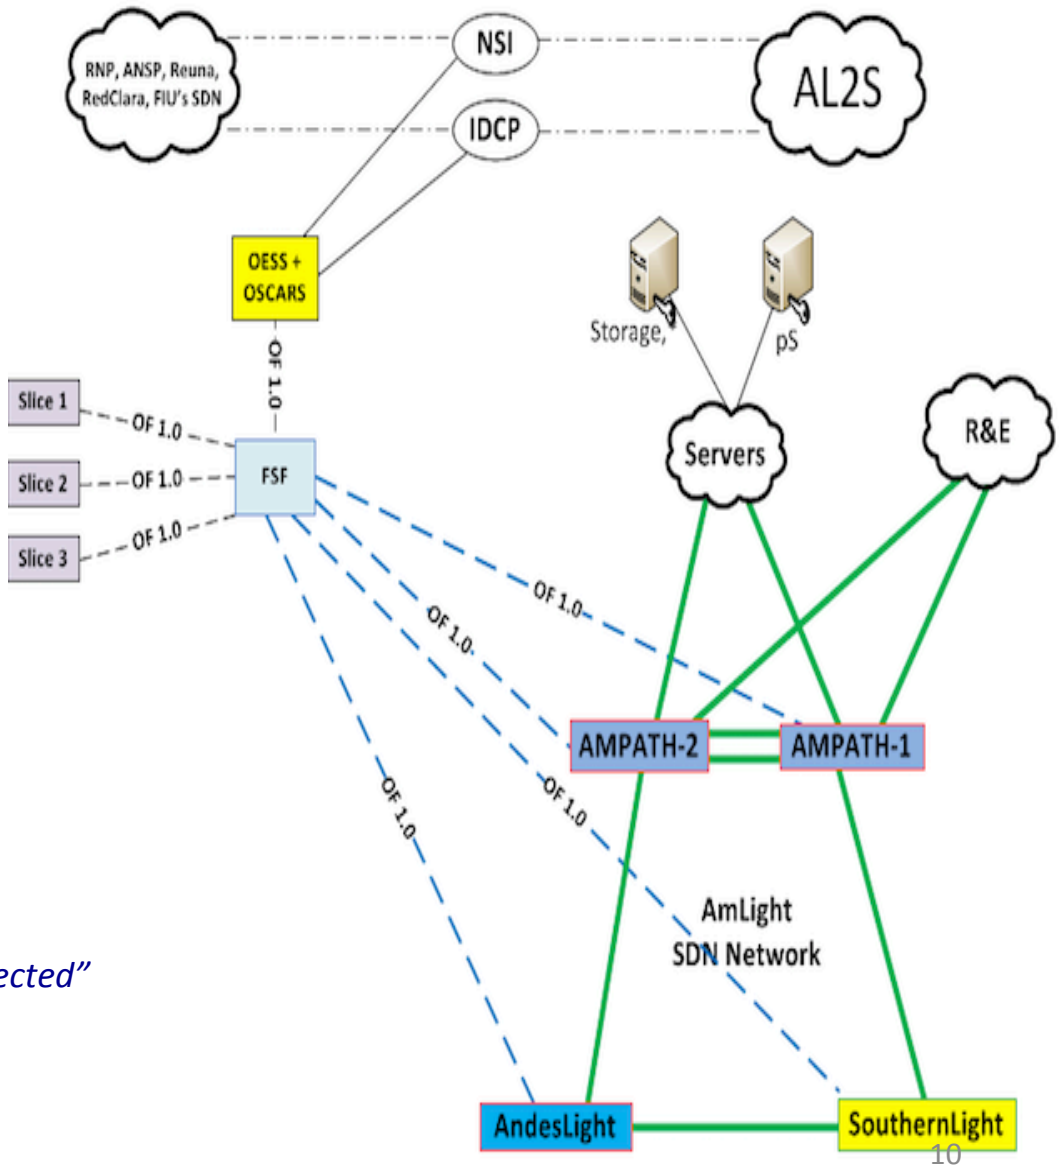

Scenario Deployed (2)

- *OESS UI:*

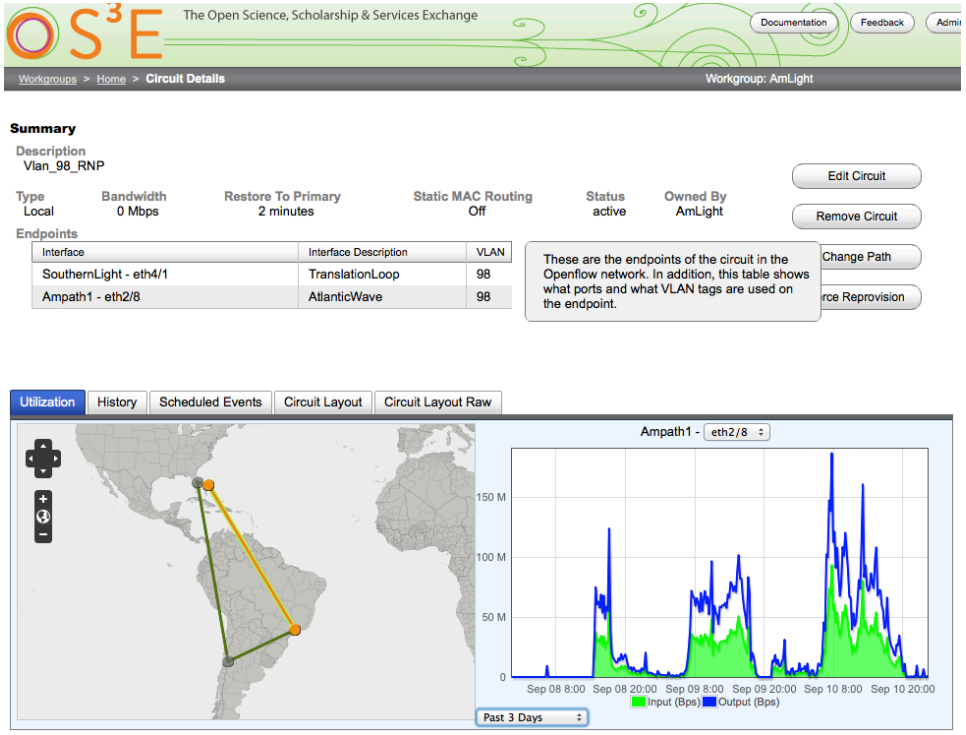

Traffic Statistics since beginning of September:

- Peak of 7,3Gbps
- OF interfaces only!

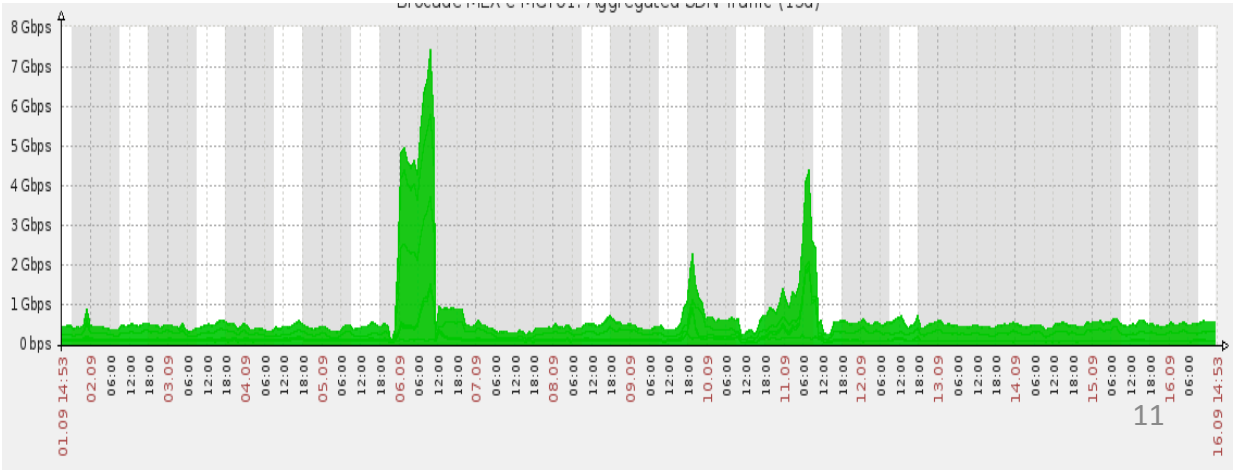

Findings (1)

A. Improving operations efficiency

	Average time to provision a new circuit		Avg. number of e-mails exchanged	
Domains involved in the path	***before SDN***	***after SDN***	***before SDN***	***after SDN***
RNP, ANSP, RedCLARA, AmLight, Internet2, ESnet	5 days	< 5 minutes	10	0
Other domains using OSCARS or NSI support	12 days	< 5 minutes	65	0
Other domains not using OSCARS or NSI support, < 3 networks in the path	5 days	*	10	*
Other domains not using OSCARS or NSI support, >3 networks in the path	12 days	*	65	*
With domains in other continents not using OSCARS or NSI support	45 days	*	100	*

* - Not tested yet

Findings (2)

B. Introducing network programmability

	Network Access and Programmability	
	Before SDN	After SDN
Network View	SNMP	SNMP and Openflow
Provisioning Defined by the User	-	Full Openflow access through a dedicated slice
Multipath experiments	Static paths offered	
Flow controlled hop-by-hop	-	

Network programmability is the main achievement of this project:

- *Network-aware applications will have AmLight as a real platform for innovation*

Some Lessons Learned

- Some legacy protocols and old switching line cards could increase the complexity
  - LACP, Counters, Ethertypes
- Lack of the Hybrid port feature can derail the SDN deployment
- Out-of-band/Control Plane network could be challenging
- “Having a testing environment with the same devices is mandatory”
- Convergence methodology has to be improved
  - Specially in long-haul links

Future

- Explore and add new features related to troubleshooting and security
- Create a *Software-Defined Internet Exchange* (SDX) involving AmLight and AtlanticWave
- Find researchers interested in using the AmLight infrastructure
- Migrate to Openflow 1.3
  - Metering and improve the network convergence

Moving Towards SDN @ AmLight

Visit our website: www.sdn.amlight.net !

Thanks!
Questions?

Jeronimo Bezerra <jbezerra@fiu.edu>