



NANOG 63
February 04th 2015

Migrating AmLight from legacy to SDN: Challenges, Results and Next Steps

Jeronimo A. Bezerra
Florida International University
<jbezerra@fiu.edu>



Agenda

- Who are we?
- What is SDN?
- Step by Step with use case
- Future

Who we are

- AMPATH:
 - Academic IXP based in Miami
 - Interconnects Latin America RENS to other RENS in the world
- AmLight:
 - A set of 4 x 10G links used to connect Latin America RENS to AMPATH
- Partnership among FIU, NSF, ANSP, RNP, RedClara, AURA and REUNA

Recap: What is SDN?

- *Software Defined Network* **decouples** Control Plane from Data Plane:
 - Forwarding decision managed by an external network **controller**
- Standard interface for the communication between the network controller and the switches:
 - A **standard** protocol was developed: **Openflow**
 - Openflow 1.0 and 1.3 is deployed and supported by lots of vendors
 - Most of the current switches which supports OpenFlow supports Hybrid Mode
 - Some ports using OpenFlow, some ports using legacy protocol
 - Some switches also support **Hybrid Ports**
 - OpenFlow and legacy traffic in the same port
 - Useful for an easy and incremental deployment

But what is SDN? (2/2)

- With SDN, the network controller would be responsible for **all** network configurations:
 - Network connectivity, including a **loop-free topology**
 - Rate-limits, **priorization**, statistics
 - And **new** services/deployments:
 - Security, new protocols, new **applications**, etc.
- Do you want to go deeper?
 - Open Networking Foundation:
 - <https://www.opennetworking.org>
 - “Software-Defined Networking: The New Norm for Networks”
 - “SDN Migration Considerations and Use Cases”
 - Openflow Specifications:
 - <https://www.opennetworking.org/openflow-conformance-certification>



Step by Step to migrate to SDN: AmLight Use Case

Step 1: Know your network

- Document everything you have in operation:
 - Link Aggregation/LACP,
 - VLANs/Spanning-Tree/QinQ,
 - MPLS/Routing/L2VPN, L3VPN, QoS,
 - Fast ReRoute, BFD,
 - Port Mirroring, IPS/IDS
- Knowing what you have will help you choose controllers and applications
 - It doesn't mean you will find applications that support everything!

Step 1: Know your network (2)

- AmLight's Operations are based on Layer 2
 - Layer 3 services are managed between users and AMPATH
- At AmLight we had:
 - Link Aggregation (MCT + LACP)
 - VLANs + per VLAN Rapid Spanning-Tree
 - Port Mirroring

Step 2: Assessments of your devices

- Supported protocols:
 - Openflow
 - Which version?
 - Netconf, Yang?
- Openflow implementation phase:
 - Beta, Testing or stable?
- Openflow's *Optional* features
 - Metering, Port Group, LACP, etc.
- Is Hybrid port supported?

Step 2: Assessments of your devices (2)

AmLight use case:

- Brocade MLXe/XMR/CES switches:
 - Openflow 1.0 and 1.3 (1.3 started on Dec 2014)
 - Amount of flows supported
 - MLXe (-D)/XMR: 64k flows (per-system and per-module), 4k per port
 - CES: 4k flows (L2) or 2k flows (L2/L3)
 - Support for Hybrid port
 - MLXe/XMR: 2k Protected vs 4k Unprotected VLAN IDs
 - CES: Doesn't support Hybrid Port
 - Amount of controllers supported
 - MLXe/XMR and CES: 3 (active or passive)
 - SSL optional (max of two)
 - Kind of matches supported:
 - MLXe/XMR: L2, L3, L2/L3 (L2/L3 only on 8x10G and 2x100G)
 - CES: L2 and L2/L3 (L3 in future)
 - Default actions:
 - Drop packets or Sent to controller
 - Statistics per Flow
 - MLXe/XMR: all
 - CES: First 2k flows

Step 2: Assessments of your devices (3)

- Brocade MLX/XMR/CES (cont.)
 - Layer 2 Control Protocols not supported on Openflow ports (even hybrid ports)
 - No STP, RSTP, etc.
 - Only 40 legacy VLANs supported on Hybrid Ports
 - No switching from LOCAL and NORMAL ports
 - Openflow and Legacy don't talk between them
 - Gen1 switching modules (4x10G) only support LLDP for generic flows
 - Also have counters limitations
 - Link Aggregation is not supported
- Network devices' hardware will always matters!
 - There will always be limitations due to the chosen internal components

Step 3: Openflow Controller and Orchestrator

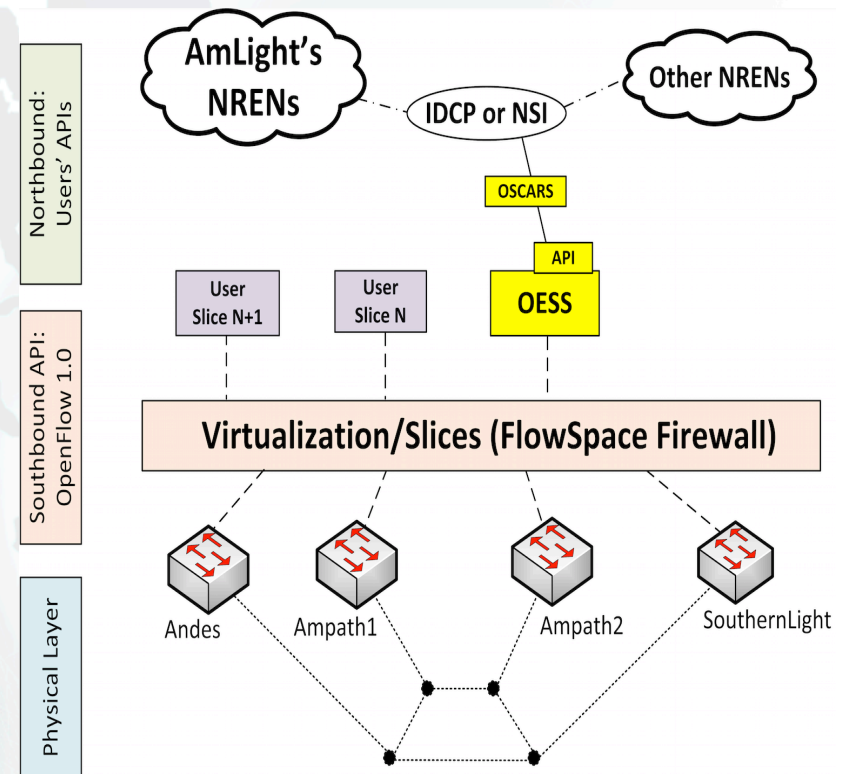
- Controller vs Orchestrator
 - Controller: manages the southbound interface
 - Orchestrator: business application
- In-house development? Use one available?
 - Does it support your applications/services in use?
- Do you need network virtualization/slicing?

Controller	Interface(s)	Lang	Modules				Licens	OF 1.3
			Dissemin	Host	Transp	MultiProc		
NOX	Native/WS	C++	Yes	yes	yes	no	Apach e	branch
POX	Native/WS	Pytho n	Yes	Yes	Yes	no	GPLv3	
FloodLight	Native,WS	Java	Yes	Yes	Yes	no	Apach e	no date
OpenDayLigh t	Native,WS	Java	yes	yes	yes	yes	EPL	yes
RYU	Native/WS	Pytho n	yes	yes	Yes	Yes	Apach e	yes

Step 3: Openflow Controller and Orchestrator (2)

AmLight Use Case:

- OpenFlow 1.0
- Controller: NOX
- Orchestrator: Internet2 OESS
 - Supports Layer 2 provisioning via Web User Interface
 - Supports OSCARS (multi-domain provisioning – useful for RENs)
- New feature added:
 - Network Virtualization: FlowSpace Firewall



Step 4: Openflow control plane network

- Where to place the controller?
 - Important question for WAN, not that much for Campus/Datacenters
 - How many controllers?
 - One per site, One per domain?
- How to reach all network devices from controller?
 - In band?
 - Most vendors don't support Openflow messages over Openflow flow entries
 - Out of Band?
 - Is there connectivity restraint? New interfaces required?

Step 4: Openflow control plane network (2)

AmLight Use Case:

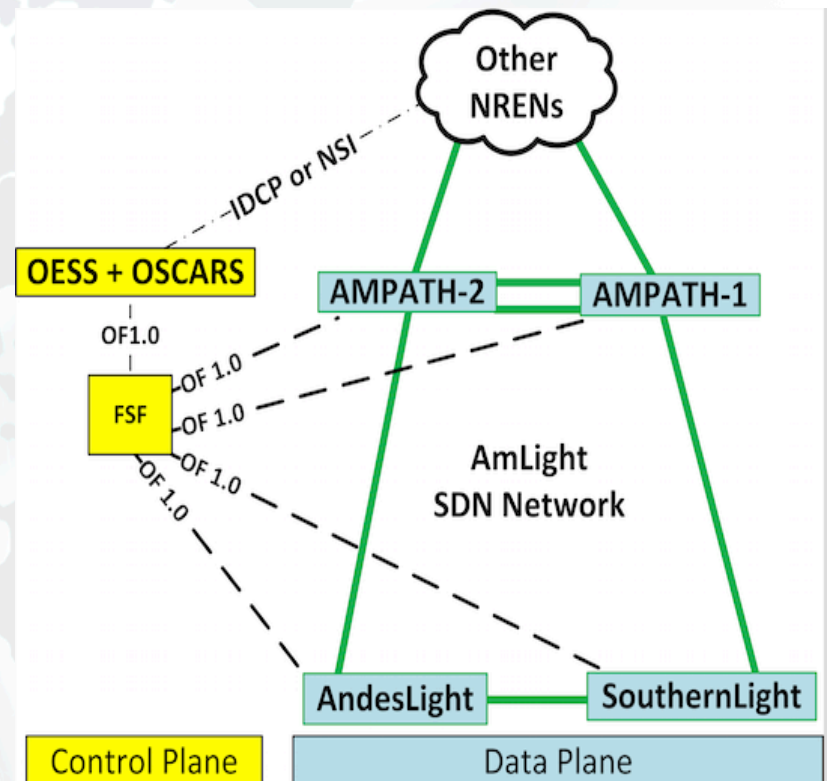
- 2 switches in Miami
- 1 switch in Brazil (SouthernLight)
- 1 switch in Chile (AndesLight)

Ring topology:

- Just two links per switch
- Chile's switch doesn't support hybrid port

Final configuration:

Out of band through a third party network
Controller in Miami (closed to the "SDN" Engineers)



Step 5: Security

- Use SSL or not for the control plane?
 - With SSL: secure communication, not all controllers support. Hard to troubleshoot (*tcpdump*)
 - Without SSL: insecure, all controllers support, easy troubleshooting
- Control the **amount of flows** per slice/virtualized network
- Create **flow insertion rate-limit** per slice/virtualized network
- How to troubleshoot in a passive way through a secure approach?

Step 5: Security (2)

AmLight Use Case:

- Use SSL or not for the control plane?
 - Without SSL: easy troubleshooting
 - Limited to 4000 L2/L3 flows (360 in use)
 - 40 flows/sec flow insertion rate-limit (15 observed)
- How to troubleshoot in a passive way through a secure approach?
 - Working on it!

Step 6: Deployment

- New skills required: Linux, Log Reading, **Coding** (Python or Java), etc.
- Start with *mininet*, try to reproduce your services and configurations
- Create a **testing** environment with real switches and, if possible, the same vendors and models
- If possible, deploy it **gradually**, for example, start with layer 2 services

Step 6: Deployment (2)

AmLight Use Case:

- Started on April 30th, Deployed on August 30th
- A few tools developed for troubleshooting
- New skills acquired: Openflow and Python (Linux was part of our routine)

Lessons Learned

- Train your team, document everything
 - Be prepared to develop some code: no controller will support everything you need
- Keep your testing environment available:
 - Every required change and/or upgrade in the future must be tested/applied on this environment
- Be prepared for surprises and crashes:
 - No vendor/switch is mistake-free.
 - Have your Disaster Recovery Plan ready

Future

- Define a methodology to
 - test new vendor/controller software;
 - validate new Openflow switches;
 - troubleshoot;
 - validate the OpenFlow forwarding actions are working;
- Deploy New Applications?
 - DDoS control (ONS 2014 IDOL)
 - Software Defined Internet Exchanges (SDX)

AmLight Today: Results

Provisioning:

<i>Domains involved in the path</i>	Average time to provision a new circuit		Avg. number of e-mails exchanged	
	<i>before SDN</i>	<i>after SDN</i>	<i>before SDN</i>	<i>after SDN</i>
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

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	-	



NANOG 63
February 04th 2015

Questions? Comments?

www.sdn.amlight.net

Migrating AmLight from legacy to SDN: Challenges, Results and Next Steps

Jeronimo A. Bezerra
Florida International University
<jbezerra@fiu.edu>

