Coexisting Production and Experimental Testbeds: The AmLight experience after one year

Jeronimo Bezerra, Julio Ibarra and Heidi Morgan
Florida International University
Miami, Florida
{jbezerra, julio, heidi}@fiu.edu
+1 305 348 4105

Marcos Schwarz and Humberto Galiza
Rede Nacional de Ensino e Pesquisa
Rio de Janeiro, Brazil
{marcos.schwarz, humberto.galiza}@rnp.br
+55 21 2102 9660

Keywords - Software-Defined Networks; research and education network; OpenFlow; slicing; network testbeds; network-aware applications.

ABSTRACT

Network-aware research applications increasingly require the capability to program network functions in order to satisfy particular requirements, such as end-to-end visibility, etc. For the successful verification of network-aware research applications, the construction and operation of experimental testbeds over international research and education (R&E) networks is of strategic importance. Supporting experimental testbeds over production international R&E networks exposes the production environment to potential risks, and increased complexity. To mitigate risks operating experimental testbeds, the AmLight project implemented a validation process, then developed and deployed the Testbed Sanitizer. Our experiences and findings using both the validation process and the Testbed Sanitizer and future work towards the support of experimental testbeds in a production R&E environment are described.

I. INTRODUCTION

When AmLight [1] deployed OpenFlow and became a production Software Defined Network in mid-2014, programmability emerged as a new service offered to the academic community. For the first time, U.S. and Latin American researchers could use AmLight's infrastructure to prototype their network-aware applications, using an approach called *slicing* [2]. With dedicated network slices, network testbeds are able to control how packets are forwarded on a per-hop basis, and if desired, packets could be manipulated along the forwarding path. Researchers now have the possibility of implementing a testbed with real network devices, including their limitations. Their applications can be moved from a simulated environment to an at-scale experimental environment, where they would face production challenges, such as CPU, memory and bandwidth limitations.

Running experimental testbeds in a production network environment involves risks and, complex operation and troubleshooting processes. Even though AmLight is a production R&E network infrastructure, network testbeds have been widely supported, and researchers encouraged to use AmLight for prototyping applications. Since 2014, more than seven testbeds were created in parallel with production applications. Each experimental testbed installed so far is completely different from every other, not just in its usage and purpose, but also in its requirements for deployment and testing methodology. To lower the risks of unexpected downtimes caused by experimental testbeds, a validation process was incorporated to confirm that OpenFlow messages coming from testbeds would not adversely affect the stability of AmLight network devices. As a complement to the validation process, the Testbed Sanitizer was created by the AmLight team to handle undetected anomalies using a software layer that filters messages that are not supported by the network devices.

Based on the experiences obtained from all hosted network testbeds, this paper was written to share AmLight's experience in supporting production and experimental testbeds in a production high-speed network infrastructure, that carries academic and non-academic traffic. The rest of the paper is organized as follows. Section II describes how network programmability and slicing were deployed at AmLight. Section III discusses risks of supporting testbeds. The validation process and the Testbed Sanitizer are described in sections IV and V, respectively. Findings, Future work and Conclusions are described in sections VI, VII and VIII, respectively.

II. NETWORK PROGRAMMABILITY AND SLICING

Network programmability was deployed at the AmLight network using Internet2's Flow Space Firewall [3] - an OpenFlow proxy that controls what OpenFlow controllers can do on the OpenFlow switches. It makes possible a new service called "*network slicing*" or just "*slicing*" with specific ports and VLAN ranges. A network slice allows multiple tenants to share the same physical infrastructure. A tenant can be a customer requiring an isolated network slice or an experimenter who wants to control and manage some specific traffic from a subset of endpoints. With slices, a controller in one slice cannot interfere with other slices; for example, it cannot remove flow entries or overlap them.

III. RISKS INVOLVED IN SUPPORTING TESTBEDS

Risks result from code instability in the OpenFlow agents deployed on the network devices, as well as the complexity of testing all functions supported by OpenFlow controllers and applications. Furthermore, code failures might crash not just the OpenFlow agent, but the entire network device, causing unexpected network outages that could potentially affect all services. These crashes might be caused by malformed or unsupported OpenFlow messages, or even buffer/synchronization flaws embedded in the network device's operating system code. AmLight network engineers have observed network outages due to a wide range of reasons, but mostly from software failures on network devices. Because of potential risks, AmLight is more cautious about deploying new applications and testbeds in the production environment.

In most outage situations, troubleshooting proved to be difficult, because of the complex architecture involved: the SDN applications, the network slicer, and the network devices; each could cause different kinds of problems. Not all of them have good support for troubleshooting. In every situation, event logs and packet inspection was used. Unfortunately, some tools did not provide enough information, making it impossible to reproduce an event, which forced AmLight to create or customize tools to improve its troubleshooting procedures.

IV. The Validation Process

AmLight adapted the validation process in use at Internet2 [4], because of similarities in the network environment and software. Internet2 was the first network to support SDN in a production environment. The validation process is as follows: (1) Researcher fills out a form, providing details about its application; (2) Researcher demonstrates its application working using AmLight Mininet topology; (3) AmLight tests the code in its testing environment; (4) Once the application passes the validation process, a production slice is created; (5) The application is then moved to the production environment.

Once the application is running in the production slice, both AmLight engineers and researchers manage it in a joint effort. If the researcher needs to change the application - adding features or fixing its behavior - the validation process must start from the beginning. The validation process usually consumes a few weeks, due to its complexity and the testing of all requested code changes, in order for the application to work properly.

V. The Testbed Sanitizer

To enhance AmLight SDN security, AmLight network engineers developed the Testbed Sanitizer. This tool has two main purposes: (a) filtering all undesired OpenFlow messages per network device's line card and software version; and (b) logging flows based on profiles, which could be defined per match, action, switch and/or using multiple combinations.

The Testbed Sanitizer was presented at the GLIF Meeting 2015 and the Internet2 Technology Exchange 2015 [5]. Figure 1 shows the Testbed Sanitizer placed between the network slicer and the SDN application/testbeds. With this solution, AmLight engineers are able to better monitor, understand and troubleshoot the SDN environment, improving network resilience.

VI. Findings

We observed that most of the availability threatening issues were stateful, and not stateless. Stateful issues occur as a result of multiple messages, or a sequence of messages; for example, multiple OpenFlow StatReq messages in a short interval, or a combination of some specific messages with some CPU utilization in a very specific software version. As a proof of concept, the Testbed Sanitizer proved to be an interesting approach, but it would require significant development effort to address all stateful issues. Our experience made clear that AmLight should work even more closely with the network device's vendor to improve its embedded OpenFlow agent, instead of creating our own external security filters.

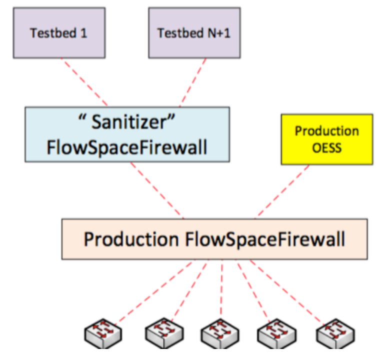

Figure 1 - Testbed Sanitizer implementation at AmLight

VII. Future Work

To improve security and ease the troubleshooting of its multi-sliced environment, AmLight is working on two other tools: an OpenFlow sniffer [6] with troubleshooting extensions and a multi-slice SDN Traceroute. The sniffer will be used to keep track of all OpenFlow messages and to help vendors and researchers reproduce and improve their code; the multi-slice SDN Traceroute is being developed to validate the data plane and guarantee consistency between all SDN layers.

VIII. Conclusions

Hosting experimental and production testbeds in parallel has proven to be complex. It requires a deep understanding of how network devices and protocols work, in order to debug issues, and to avoid impacts to network resilience. Troubleshooting tools and OpenFlow agents still need to evolve to protect against experimental applications compromising overall availability of the production environment. The validation process and the Testbed Sanitizer have considerably improved the security and the efficiency of testbed deployments at AmLight.

IX. Acknowledgment

This material is based upon work supported by the National Science Foundation under Grant No. ACI-0963053, IRNC-ProNet: Americas Lightpaths: "Increasing the Rate of Discovery and Enhancing Education across the Americas". In addition, the authors would like to thank Florida International University, Rede Nacional de Ensino e Pesquisa (RNP), Academic Network of São Paulo (ANSP), RedCLARA, Red Universitaria Nacional (REUNA), Florida LambdaRail, Internet2 and Indiana University/GlobalNOC for their support of the AmLight project.

X. References

[1] Ibarra, J.; Bezerra, J.; Alvarez, H.; Cox, D.; Stanton, M.; Machado, I.; Grizendi, E., Lopez, L.: "Benefits brought by the use of OpenFlow/SDN in the AmLight intercontinental research and education network" ,IM 2015 - 14th IFIP/IEEE Symposium on Integrated Network and Service Management, Ottawa, Canada, 11-15 May, 2015

[2] Network slice with Flowvisor, http://onlab.us/flowvisor.html

[3] FlowSpace Firewall: http://globalnoc.iu.edu/sdn/fsfw.html

[4] Internet2 Technology Exchange Presentation: https://meetings.internet2.edu/2015-technology-exchange/detail/10003939/

[5] Internet2 Technology Exchange Presentation: https://meetings.internet2.edu/2015-technology-exchange/detail/10003942/

[6] OpenFlow Sniffer v0.1, http://github.com/jab1982/ofp_sniffer

AUTHOR BIOGRAPHIES

Jeronimo A. Bezerra, MSc in Mechatronics, is the Senior Network Engineer at the Florida International University (FIU) responsible for the AmLight project and its SDN deployment and operation. He is involved in projects and experimentation related to Software-Defined Networks and also involved with the Global Lambda Integrated Facility (GLIF) and with the AutoGOLE Project.

Julio E. Ibarra, PhD, is the Assistant Vice President for Technology Augmented Research at Florida International University (FIU). Dr. Ibarra is responsible for furthering the mission of the Center for Internet Augmented Research and Assessment (CIARA) – to contribute to the pace and the quality of research at FIU through the application of advanced Cyberinfrastructure. He is responsible for strategic planning and development of advanced research networking services, including the development and management of the AMPATH International Exchange Point for Research and Education networks. He is the Principal Investigator of multiple NSF International Research Networks Connection funded projects (ACI-0963053; ACI-1451018; ACI-1451024), involving the development and operations of high-throughput international network connections to enhance U.S. e-science initiatives in Latin America, and the Caribbean.

Heidi L. Morgan, PhD. is the Director of the Center for Internet Augmented Research and Assessment (CIARA) at Florida International University. She is a Co-PI for several NSF funded projects including SwitchOn - Exploring and Strengthening US-Brazil Collaborations in Future Internet Research (switchon.ampath.net), PIRE: Training and Workshops in Data Intensive Computing Using The Open Science Data Cloud (www.opensciencedatacloud.org), Americas Lightpaths: Increasing the Rate of Discovery and Enhancing Education across the Americas (amlight.net) and the AMPATH International Exchange Point in Miami. Heidi enjoys working to advance research and education networking initiatives in the Caribbean, Mexico, Central and South America and collaborating with likeminded professionals in the US and around the world.

Marcos F. Schwarz, MSc in Computer Engineering, is R&D Coordinator at the Brazilian National Research and Education Network (RNP). He is involved with research and experimentation in Software-Defined Networks and coordinates projects related to network performance monitoring and Dynamic Circuit Networking (DCN).

Humberto S. Galiza de Freitas, BSc, is a Senior Network Engineer of Americas Lightpaths (AmLight) (amlight.net) and of the International Exchange Point for Research & Education Networking in Miami (AMPATH) (ampath.net) teams. He has been working with Internet industry since 2003, and more recently joined AmLight/AMPATH engineering teams, where he is involved in research and experimentation of Software-Defined Networking (SDN) technologies.