Handling Network Events in a Production SDN Environment: the AmLight Use Case

Jeronimo Bezerra[1], Jorge Marcos[2], Julio Ibarra[1], Heidi Morgan[3] and Luis Lopez[2]
jbezerra@fiu.edu, jrgmrs@ansp.br, julio@fiu.edu, hlmorgan@isi.edu, lopez@ansp.br

[1]Florida International University
Miami, USA

[2]Academic Network of Sao Paulo
Sao Paulo, Brazil

[3]University of Southern California
Los Angeles, USA

Keywords - Software-Defined Networks; research and education network; OpenFlow; troubleshooting; network testbeds.

I. Introduction

A generous range of tools is available for network monitoring and measurement of legacy networks. Most network protocols are well known with a variety of open source and commercial tools to help network engineers with their daily operation. Network Monitoring Systems (NMS) can handle SNMP GETs, ping probes, and traceroute routines natively. Some SNMP MIBs were standardized by the IETF to collect protocol-specific information, such as VLAN utilization and OSPF LSAs. With OpenFlow and SDN, some network characteristics were changed, making some tools and protocols less useful or, sometimes, completely useless. In the past, network engineers assumed the control plane inside the network devices was working properly, focusing most of the troubleshooting process on the data plane or configurations.

The lack of specific OpenFlow tools and debugging protocols makes troubleshooting of SDN networks extremely more complex. In OpenFlow and SDN environments, with an entirely new control plane in place, two additional components must be included in the troubleshooting scenario: one is the OpenFlow agent inside the network device, and the second is the OpenFlow controller and its associated application. There is added complexity the moment network engineers have to troubleshoot applications. In some cases, specific knowledge of programming languages is required, a skill almost ignored by most of the network engineers.

The complexity involved in the management of OpenFlow and SDN environments has significantly affected and still affects the AmLight[1] network operation. Once AmLight was migrated to an OpenFlow-based network in 2014 [1], new skills were gained, new troubleshooting tools were incorporated, and some scripts and tools had to be created or customized. To make things even more complex, network events in the AmLight production environment must be handled in the least disruptive way possible (different from network testing environments or simulations) because the troubleshooting process cannot affect production traffic. However, sometimes, extreme actions have to be taken. Finding a balance is an art at this moment, and having the appropriate set of tools is fundamental.

In past presentations at Internet2, TNC, and GLIF meetings, AmLight presented its experience in supporting experimental testbeds in parallel with production traffic from a high-level perspective. Now, AmLight engineers aim to present how the network operation happens at AmLight, including tools developed specifically for SDN troubleshooting. As part of this presentation, we will describe some challenges when troubleshooting OpenFlow networks (Section II); what tools had to be improved and how these tools fit in the AmLight network-monitoring environment (Section III). Section IV reports future work.

II. Challenges when troubleshooting OpenFlow-based networks

Because OpenFlow is a new protocol, its implementation has to be carefully planned in production environments. In an OpenFlow-based network, the troubleshooting process includes handling the Control Plane, the Data Plane and also the Management Plane. Each plane brings specific challenges. Some of the Control Plane's challenges involve (i) understanding the impact of a new southbound protocol in the network device; (ii) how to track the southbound communication; (iii) how to store and replay such communication for future troubleshooting activities; (iv) understand the set of features properly supported by the network device; and (v) how network devices handle malformed messages. Monitoring the control plane can be done in a completely passive approach, using sniffers and logs. The Data Plane requires an active approach: the OpenFlow controller has to send network probes to validate the data path. The active approach is required to evaluate the consistency between multiple layers in the SDN network

[1] Americas Lightpaths (AmLight) is a project of the U.S. National Science Foundation International Research Network Connections (IRNC) program to facilitate science research and education between the U.S. and the nations of Latin America. AmLight is a production network composed of a number of international network links connecting U.S. R&E networks to similar networks in Latin America.

stack, from application to OpenFlow agents to network device's forwarding memories. The Management Plane offers challenges when network devices do not properly deploy it, for instance, when counters and timers advertised by network devices are not accurate, compromising the understanding of the network.

Furthermore, other troubleshooting challenges related to OpenFlow-based networks are important and ignored by the academic community: vendor implementations and support for troubleshooting inside the network devices. In the troubleshooting process, it always comes to the point when engineers will need to use the Command Line Interface (CLI) provided by the vendor to deeper investigate the issue. Without a properly implemented CLI, troubleshooting might become impossible, turning into a series of try-and-fail attempts, compromising the network resilience even further.

III. NETWORK TROUBLESHOOTING AT AMLIGHT

Figure 1 represents the SDN architecture proposed in the IETF RFC 7426 [2]. Many academic works were published focusing on specific components and layers of this architecture, for instance [3] and [4]. Not a single work became production software or created an integrated solution covering many planes. At AmLight, different SDN troubleshooting tools were developed to focus on specific problems of each of the SDN planes. The OpenFlow Sniffer [5] was created to track the communication between OpenFlow controllers and network devices. This tool was launched in 2015 as demand for a production tool capable of handling OpenFlow 1.0 in a sliced SDN environment.

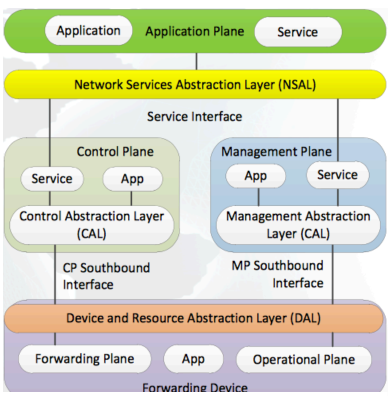

Figure 1 SDN Architecture/RFC 7426

At the same time, the AmLight's Zabbix [6] NMS was extended with a few scripts that collect and consolidate OpenFlow data, such as the number of flows, number of PacketIn, PacketOuts and FlowMods per device per second. Zabbix uses a database to store all data collected, making it possible to have historical data to correlate events. Through the scripts created, reports are provided per slice and for the overall network topology, optimizing the network troubleshooting.

In 2016, the troubleshooting focus turned to the SDN *black holes*. Black holes happen when a specific OpenFlow entry in the network device's memory (TCAM, for instance) corrupts and discards packets instead of forwarding them. From the Control Plane perspective, the flow entry is installed, sometimes counters are even increasing, but packets are not being forwarded. In a network with multiple paths, hundreds of flow entries and multiple network devices, tracking a specific black hole can be very time-consuming. The SDNTrace [7] was developed to identify the physical data path and, at the same time, find black holes in a more automated approach. The idea behind the SDNTrace was inspired by [8]. Even though some SDN applications have a data path validation, they only validate a link between switches, not user flows.

IV. FUTURE WORK

Having multiple tools in a decentralized environment is not the most optimized approach in any network configuration, SDN or legacy. Zabbix is consolidating OpenFlow data; the OpenFlow Sniffer tracks the Controller-to-Switch communication; and the SDNTrace is used for active monitoring. There is no communication or integration between these tools, mostly because they were not planned to. To overcome this lack of communication, AmLight engineers are developing a new approach to monitoring SDN networks, as a side application that will consolidate all data and active tests into a single REST and Web User Interface. The idea is to evolve how SDN monitoring and management is performed.

For the TNC 2017, AmLight engineers plan to present not just a new approach for network monitoring, but also the SDN tool being developed. This tool will support different OpenFlow versions and different OpenFlow controllers, such as Ryu, ONOS, and OpenDayLight. Also, this new tool will be launched as Open Source code.

REFERENCES

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

[2] RFC 7426 - Software-Defined Networking (SDN): Layers and Architecture Terminology, IETF, 2015

[3] Chowdhury, Shihabur Rahman, et al. "Payless: A low cost network monitoring framework for software defined networks." 2014 IEEE Network Operations and Management Symposium (NOMS). IEEE, 2014.

[4] Handigol, Nikhil, et al. "Where is the debugger for my software-defined network?." Proceedings of the first workshop on Hot topics in software defined networks. ACM, 2012.

[5] Bezerra, Jeronimo, et al. "AmLight's OpenFlow Sniffer dissected: Troubleshooting production networks." Anais do WPEIF 2016 Workshop de Pesquisa Experimental da Internet do Futuro.

[6] Zabbix NMS, http://www.zabbix.com

[7] AmLight's SDNTrace: https://github.com/amlight/sdntrace

[8] Agarwal, Kanak, et al. "SDN traceroute: Tracing SDN forwarding without changing network behavior." Proceedings of the third workshop on Hot topics in software defined networking. ACM, 2014.

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.

Jorge Marcos de Almeida, Electrical Engineer, is the network engineer at the Academic Network at Sao Paulo (ANSP), Brazil, responsible for the network configuration, deployment and operation. He is involved in projects and experiments related to Software-Defined Networks and cloud computing. He is also involved with the Global Lambda Integrated Facility (GLIF) and with the AutoGOLE Project, being responsible for the SouthernLight GLIF GOLE.

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.