Enabling Network Operating System Automated Upgrade in Packet/Optical Whiteboxes

Samier Barguil
Universidad Autónoma de Madrid, Madrid, Spain

Roman Makarov, Murat Mugan,
Facebook Connectivity,
Menlo Park, USA

Juan Pedro Fernandez-Palacios, Victor López
Telefonica I+D
Madrid, Spain

***Abstract*— This work describes and implements an automatic network operating system upgrade in a Whitebox environment. The network operating system upgrade workflow moves the traffic thorough the packet/optical whiteboxes during the maintenance window and return it after the upgrade is complete.**

Keywords—(060.4250) Networks; (060.4510) Optical Communications

I. INTRODUCTION

Nowadays, the datacenter and network service providers are in a continuous race to satisfy the variable market needs, because the demands of capacity are getting bigger but also the reliability and the continuity plans are becoming fundamental in the commercial offers. So, in order to achieve competitive KPIs as well as competitive prices the network infrastructure deployed becomes an important part of the business. Historically, networking components have been managed by a close-set of reliable set of providers. However, since several projects have been released to open the market of those components to new competitors. In 2011 the Open Compute Project (OCP) [1] started with the objective to create of open hardware specification (servers and racks) for datacenter usage and in 2016 the Telecom Infra Project (TIP) [2] was released with the aim of provide open hardware for telecommunications. Those open initiatives have led to the reduction of CapEx investments in infrastructure as well as some OpEx reductions based on the optimization of the power consumption or the selection of the "base" components.

In the majority of the cases, the open-hardware inclusion in the telecommunications infrastructure is focus on the disaggregation of some of the functionalities, including switching (rack interconnection) and transmission infrastructure (WAN interconnection). Disaggregation in a device is based on the separation of the software known as the Network Operating System (NOS) and the hardware that becomes a whitebox as it can be used with any NOS [3]. The disaggregated scenario allows the operation and integration of a real multivendor scenario, where a new NOS can be deployed in the same whitebox that is in production. This avoids the common problems of time to market as the hardware migration is not required. The hardware used for these disaggregated scenarios varies according the technology, however for transport networks, TIP Voyager and Cassini are some examples of the leading trend. They are open designs of optical transponders with packet switching functions used for WAN interconnections.

As we mention before, the adoption of this kind of hardware has another set of advantages more related with their integration possibilities. Due to the hardware and software decoupling, the hardware can run different NOS appliances. Enabling with this the testing of automatic solutions of common tasks and the reduction of the human intervention in the day-to-day operations in the datacenter. This was hard to achieve in the past, because the entire configuration done in equipment were vendor-dependent and each manufacturer implements their proprietary interface. Sometimes, those interfaces were

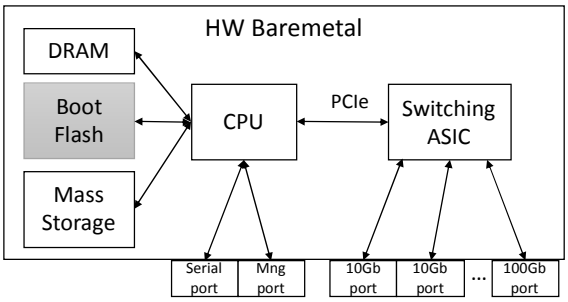

Fig. 1. HW resources available with ONIE

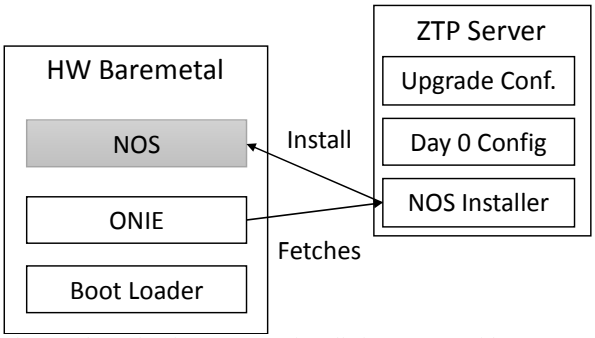

Fig. 2. Schematic of remote NOS installation process with ONIE

978-3-903176-21-8 © 2020 IFIP

different from software version-to-version or even from hardware model to model.

In recent years several automation initiatives have been released to automate common tasks in the datacenter environment, Facebook's Robotron [4] and Google's Zero Touch Network [5] are frameworks designed to minimize human errors in management tasks. There are works on the automation of service provisioning in disaggregated optical networks [6,7], but there are no scenarios how to automate the NOS operation in production environments with open whiteboxes. This work describes and implements an automatic network operating system upgrade in a whitebox environment.

II. Network Operating System Provisioning and Operation

Central offices and data centers are complex and controlled environments were multiple technologies coexist. In those environments Networking. Computing and Storage devices are continually manipulated and operated following strict protocols to fulfill high reliability and business continuity standards. One of the most common and critical changes done in those environments is the software upgrade. Software upgrade process in servers or applications is a daily task that is managed in a seamless fashion by the software providers. As an example, any smartphone upgrades applications periodically without troubles for the end-user. Network infrastructure is becoming more a more a software consumer. However, the software upgrade in a router or a transponder is a tedious process that requires a very specific procedure that is vendor oriented.

White box solutions have a common procedure to install a NOS thanks to the utilization of Open Network Installation Environment (ONIE). ONIE is preinstall in most of the whiteboxes and allows the utilization of the CPU, memories (RAM, boot flash and mass storage) and serial and management ports of the device (Fig. 1). It does not provide any capability to use enables the data plane ports via the switching ASIC. However, ONIE enables local installation of a NOS from a USB or downloading an image from a remote server (FTP, HTTP). This option to remotely look for a NOS image enables the implementation of a Zero-Touch Provisioning (ZTP) solution. ZTP is the process to deploy a NOS and a base configuration in a network element, thus is a router or a transponder, so the network element can enter in production without any human configuration. ZTP process is done for the first time the device is turned on in the network. Periodically, system vendors release new versions of their NOS. A very similar process can be done for this upgrade scenario. Nowadays, the NOS upgrade process is a vendor dependent process and the procedure differs between each vendor solution. The ZTP and upgrade process presented in this part is a first step to have a common procedure for open white box scenarios.

III. Workflow

ZTP process is based on ONIE capabilities. When installing a Cassini box, the device is running only with ONIE. At startup phase, ONIE requests for an IP address via DHCP, which is solved by a ZTP server. ZTP server answers with the URL of the NOS installer, so Cassini can download the image. URL can be an exact URL or a partial URL that contains variables adapted for the operational area. NOS image is downloaded to the device and it is installed. Once the installation is complete, the device applies the day-0 configuration. To do so, it requires a new dialogue with the ZTP server, which may include a license activation process.

The NOS upgrade process is similar to ZTP with the main difference that there are services in production. Fig. 3 represents the NOS upgrade workflow. This process will be done in a planned maintenance window. At the initial state, the network is running and operational with services provisioned via the Cassini devices. End user applications are moved to the backup path (gracefully) via IGP protocol. At this stage link-using Cassinis don't carry any data plane traffic. We carry out NOS upgrade process in the Cassini(s) using ZTP via management interface. ZTP process includes firmware download, config load and license install using DHCP options. Once the upgrade process is finished, the traffic moves back to primary link.

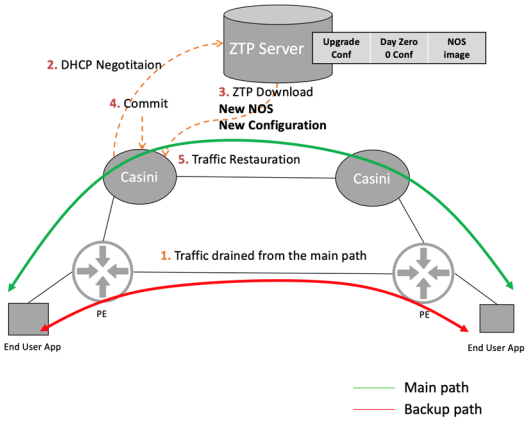

Fig. 3. Steps of the ZTP process

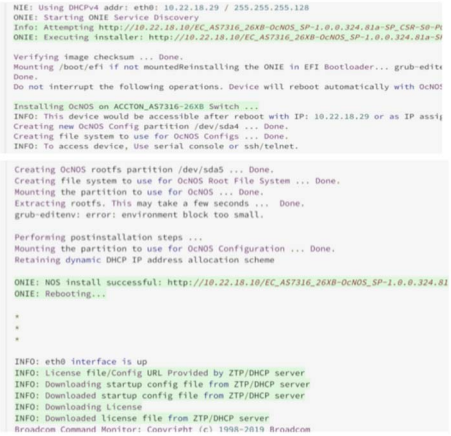

Fig. 4. Zero Touch Provisioning results from the Casini's logs

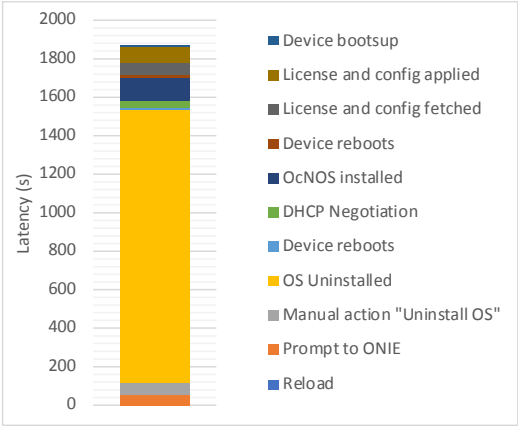

Fig. 4. Latency for the ZTP process

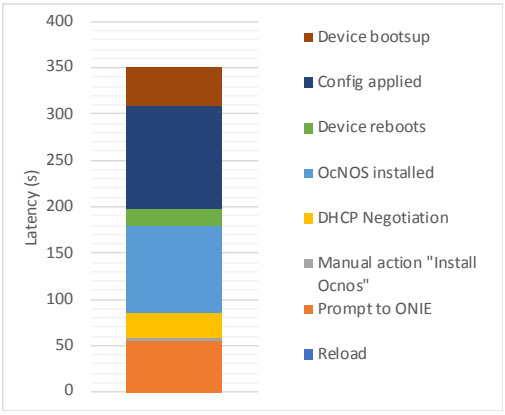

Fig. 5. Latency for the NOS upgrade process

IV. Experimental Validation

The experimental setup was deployed at Facebook facilities in Menlo Park. The IP layer was based on two routers acting as PEs towards the end user applications. For sake of simplicity they were emulated as pings, authors in [8] work on hitless distributed video applications; using programmable switching hardware to dynamically filter and forward the traffic stream to the correct paths. The IP/transport network is based on two Edgecore Cassinis. Cassini whiteboxes may act as PEs, but they are based on Broadcom's StrataXGS Tomahawk, which does not provide all IP capabilities for service provider environments. Fig. 4 shows the logs of the ZTP workflow from the Casini's logs where all the installations, license application and target configuration are executed.

Fig. 5 shows the latency of each step in the ZTP process. Most of the time is consumed by the uninstallation ("OS Uninstall") initial process, which consumed 814s, thus is 65,9% of the total time. The other two tasks that require more time are "OcNOS installed" (120s – 9,7%) and "License and config applied" (107s – 8,7%). DHCP negotiation requires only 32s of the entire process. "OcNOS installation" process includes the image download, license and the configuration application (Logs depicted in in Fig 4.). Data center and central office environments have high-speed connections, but other scenarios like Mobile backhaul will have different number as the connectivity may not have such high bandwidth connections. Fig. 6 displays the average latency results for the NOS upgrade process; the more important steps "OcNOS installed" (93s – 26,4%) and "Config applied" (111,3s – 31,7%)

V. Conclusions

NOS upgrade is not a seamless process in network operator scenarios. The advent of open white boxes enable to rethink how the NOS management (ZTP and upgrades) can be done. This article motivates the need to define an automated procedure to enable NOS upgrade in WAN networks. Moreover, there is an implementation of the workflow to carry out a ZTP process and upgrade a commercial NOS and Cassini white boxes. The NOS upgrade workflow routes the services to a secondary path and during the maintenance window, carries out the upgrade and returns the services after the upgrade is complete. The results have shown the ZTP procedure is scalable and highly desirable to implement in controlled environments such as datacenters, because with a common infrastructure and using well-known protocols such as DHCP it is possible to manage the software versioning of the networking devices. The workflow used enables the automation of the software upgrade process and allows the control of the current and target configurations. The current workflow must be adjusted for scenarios where the bandwidth is limited or the management connection is not independent from the network connection; in those cases the time expend in the file transmission and the network reliability are the main variables to evaluate.

References

[1] Open Compute Project, https://www.opencompute.org/

[2] Telecom Infra Project, https://telecominfraproject.com

[3] V. López, O. Gonzalez, J.P. Fernandez-Palacios: Whitebox Flavors in Carrier Networks, in Optical Fiber Conference (OFC), March 2019.

[4] Y.W.E. Sung et al., "Robotron: Top-down Network Management at Facebook Scale", in Proceedings of the ACM SIGCOMM, 2016.

[5] B. Koley, "The Zero Touch Network", in Proceedings of IEEE International Conference on Network and Service Management (CNSM), 2016.

[6] R. Morro, et al.: Automated end to end carrier ethernet provisioning over a disaggregated WDM metro network with a hierarchical SDN control and monitoring platform. Proc ECOC 2018.

[7] A. Mayoral Lopez-de-Lerma, et al.: Multi-layer service provisioning over resilient Software-Defined partially disaggregated networks, in Proc. OFC, Mar 2019.

[8] B. Andrus, B. M., Sasu, S. A., Szyrkowiec, T., Autenrieth, A., Chamania, M., Fischer, J. K., & Rasp, S. (2019, March). Zero-Touch Provisioning of Distributed Video Analytics in a Software-Defined Metro-Haul Network with P4 Processing. In 2019 Optical Fiber Communications Conference and Exhibition (OFC) (pp. 1-3). IEEE.