The Need for SDN in Orchestration of IP over Optical Multi-Vendor Networks

Ori Gerstel[(1)], Victor Lopez[(2)]

[(1)] Sedona Systems, 17 Hatidhar St. Raanana Israel; email: orig@sedonasys.com

[(2)] Telefónica I+D/GCTO, Distrito Telefonica, c/Ronda de la Comunicación s/n, 28050 Madrid, Spain

Abstract *We explain why distributed control lacks the ability to optimally control networks with multiple transport domains or both IP and optical layers. We then propose a practical architecture to fix these issues and experimentally demonstrate it over commercial IP/transport gear.*

Introduction

The merits of close collaboration with the IP and transport layers has been discussed extensively over the last few years[1-3]. Such technology can significantly increase the efficiency of the network and reduce its cost[4], improve the coordination between the layers during maintenance and optimization processes, improve SLAs, and simplify the operation of the network. In the past, the coordination between the layers was naturally assumed to be done with distributed control, using extensions to existing protocols like GMPLS (driving for the definition of a UNI protocol[5]), or PCEP[6]. However, the recent push for centralized control, under the umbrella of software defined networking (SDN), has caused a rethinking of the architecture.

In this tutorial paper we will examine the need for central control in the context of realistic service provider (SP) networks that comprise of multiple administrative domains, multiple layers – in particular optical and IP layers – and multiple vendors at all layers. We will show that without such centralized control, the system does not have enough knowledge to add new links in the IP layer based on traffic changes. In addition, under certain failures, the right way to restore service involves understanding of multiple domains. Even simpler tasks, like optimizing the route of optical connections cannot be done in hitless fashion without global awareness. Such global awareness is very hard to implement without central control.

It should be noted that we adopt a more near-term broad definition of SDN: a technology that exploits central control and open interfaces to simplify the control of the network and remove vendor lock in. We do not assume that network gear must be swapped out for white-boxes, or that the existing distributed control plane will be removed. In fact, we note that service providers see value in the distribute control plane. We also don't assume any specific protocol between the controller and the network elements. The value central control brings to the network is mostly independent of these details.

Multi-vendor / multi-domain optical networks

Typical SP networks contain more than one optical vendor. Since each optical network is a closed system (at least when it comes to DWDM networks, in which multi-vendor interop in the analog domain does not exist), this implies that a single vendor subnetwork forms as least one domain. Setting up a connection across multiple such domains requires understanding of the constraints in each domain and based on this deciding which of the domain entry and exit points should be used to set up the connection. Since optical feasibility is computed differently by every optical vendor, it is impossible for the controller of one vendor to decide what is feasible in the domain of a different vendor (unlike IP networks, in which multi-domain routing is simpler to achieve).

Therefore, collaboration between different controllers is needed to find a feasible solution. Adding to this various constraints on the routing of the connection (low latency, avoidance of certain SRLGs etc.), the problem becomes extremely complex if solved in distributed fashion. A central multi-domain controller or orchestrator that has the view of all domains and is capable of querying their respective controllers before setting up the connection is much simpler. Hierarchical PCE is a possible solution to this problem, but there is a lack on the agreement on the protocol to interface the central element (SDN controller or H-PCE).

Multi-layer networks

Controlling an IP layer on top of a multi-vendor optical layer adds another level of complication that is even harder to solve in distributed fashion. The main added complexity is that any change in the topology of the IP layer – such as adding a link to optimize the IP topology, or temporarily taking down an IP link while optimizing the optical

layer – must be carefully managed, based on understanding the entire end-to-end IP traffic in the network. Once the traffic is known, it is possible to simulate the impact of the topology change on the IP layer – both under normal conditions and under failures. This is necessary since a change in the IP topology can yield unexpected behavior, such as overload of the new link if its routing metric is too low, light load on the link if the metric is too high, or overload of another link under certain failure conditions. All this requires global knowledge that is hard to compile and disseminate in a distributed system.

This process also requires a sophisticated network simulation engine (or online planning tool[3]) that goes through various what-if scenarios. This process is CPU and memory intensive, and therefore a poor fit for the embedded controllers that run the distributed control plane in the gear.

Another reason for central control of multi-layer networks – at least in the near term – has to do with the desire to keep the human operator in the loop when such decisions are made. This is because topology changes have been always planned in manual fashion and it will be hard for operators to entirely give up manual control, at least until they gain confidence in the technology.

The role of distributed and central control

The need for central control does not imply that there is no value for the existing distributed control planes in routers and in optical gear. Certain functions are not likely to be fully centralized. For example, most SPs are reluctant to place the responsibility for restoration from network failures in the hands of a single controller that can fail, or just lose connectivity to the network during a fiber cut (which is not a rare event since the control network is often built using the resources of the network it is controlling). Restoration must also be a quick process that must not suffer from bottlenecks that are more likely to arise in a centralized control architecture. This is valid for the optical layer, but even more so for the IP layer, in which a plethora of restoration mechanisms are in place – from MPLS FRR, to IGP convergence. Other functions that are likely to stay in the distributed domain are path setup and teardown and various monitoring and alarming functions.

The likely role of central control is to drive the policy of the distributed control system. For example, to drive the routing constraints of optical connections during normal and failure conditions: what is the maximum allowable latency? Should the connection be diverse from connections supporting other IP links? Should the diversity be maintained for the restoration path as well? This way, the distributed control plane can react autonomously during the failure, but in accordance to the globally optimal policies defined by central control.

Central control will play a more active role for non-real time tasks such as network optimization, or reversion back to normal after the failure has been fixed. In these cases, it will directly orchestrate the process and ensure it takes place in a hitless fashion, when traffic conditions allow for it.

We believe that the central control architecture will be hierarchical: single vendor/layer controllers at the bottom of the pyramid, an orchestration platform connected to their northbound interfaces (NBI) in the middle, and multi-layer applications on top – as shown in Fig. 1. This is because each of these controllers will be built by different vendors to best control their gear and will not be necessarily based on the same common platform – at least in the near term, as various control platforms jostle to establish their leadership. While the controllers may remain disparate for some time, there are efforts to standardize the NBI, which we hope will bear fruit in the near term.

Experimental demonstration

We demonstrated this architecture at OFC'15, using Cisco and Juniper commercial core routers and optical gear from Infinera, Cisco, Huawei and Ciena. While Cisco and Infinera optical gear was co-located with the routers at the Telefonica lab, Ciena and Huawei optical gear was controlled remotely from their labs – see Fig. 1 for details. In this demo, we used Sedona Systems' Multi-Layer platform (MAP), which was integrated with the NBI of the controllers for the various gear (one controller per optical domain and one for all routers), and abstracted the vendor specific details toward the applications. We also used Sedona's ML visibility application (MLV) and its optimization application (MLO) to optimize the IP network. A screenshot from the demo is shown in Fig. 2. The full video capture can be found at [7].

The demonstration showed how the network reacts to an increasing load on IP links (simulated by increasing the bandwidth reservation of an MPLS tunnel), by automatically adding a new IP link between routers – instead of the trivial (but sub-optimal) solution of upgrading the capacity of the congested link.

MLO decides automatically which link to add, after it simulates the expected impact on the IP layer, using the Cisco IP layer online planning function, which is part of the controller. MLO also identifies an optical path for the new link that minimizes common failure points with other IP

links. The new optical path is selected in consultation with the controllers of the optical domains, considering the constraints of each domain. MLO then established the new IP link and its multi-domain optical path automatically, using the Cisco and Infinera optical layer controllers. This causes MPLS to reoptimize the tunnels and the congestion is removed.

Acknowledgements

The experimental part of this paper was enabled through close collaboration with Chris Liou, Abhinava Sadasivarao and Sharfuddin Syed from Infinera, Walid Wakim and Gabriele Galimberti from Cisco, the Telefonica team and the Sedona Systems development team. This work was partially sponsored by the ACINO project (EU grant 645127).

References

[1] Telefonica et al., “IP and optical convergence: use cases and technical requirements”, at tid.es website, 1/2014.

[2] AT&T, “AT&T Domain 2.0 Vision White Paper”, available at www.att.com website, 11/2013.

[3] L. Velasco et al., “In-network planning”, IEEE. Com. Mag., 52 (1), 52-60, 2014.

[4] O. Gerstel et al., “Multi-layer capacity planning for IP-optical networks”, IEEE. Com. Mag., 52 (1), 44-51, 2014.

[5] IETF standard, https://tools.ietf.org/html/rfc4208

[6] IETF standard, https://tools.ietf.org/html/rfc5440

[7] Multi-layer/multi-vendor control demonstration, OFC15: https://www.youtube.com/watch?v=pmzdsf-oHCw

[8] V. López et al., “Demonstration of SDN orchestration in optical multi-vendor scenarios”, OFC 2015.

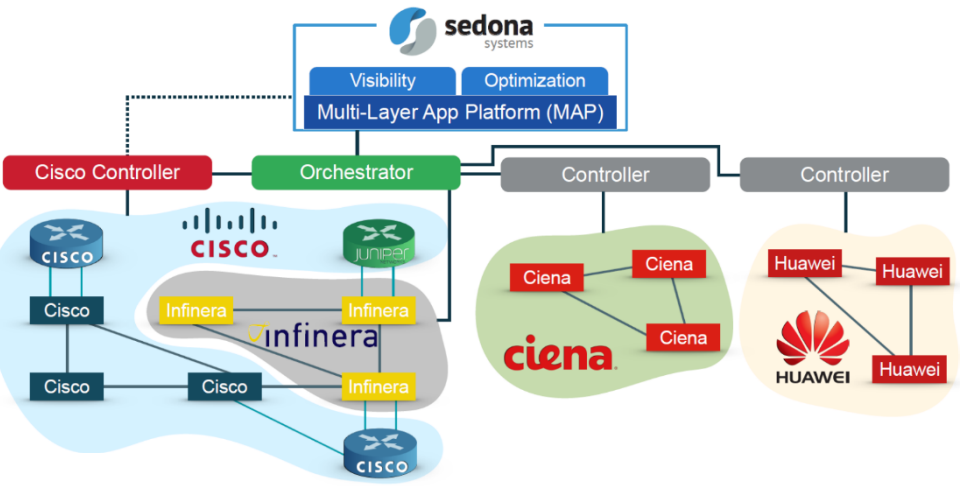

Fig. 1: Setup of the multi-vendor and multi-domain demonstration at OFC

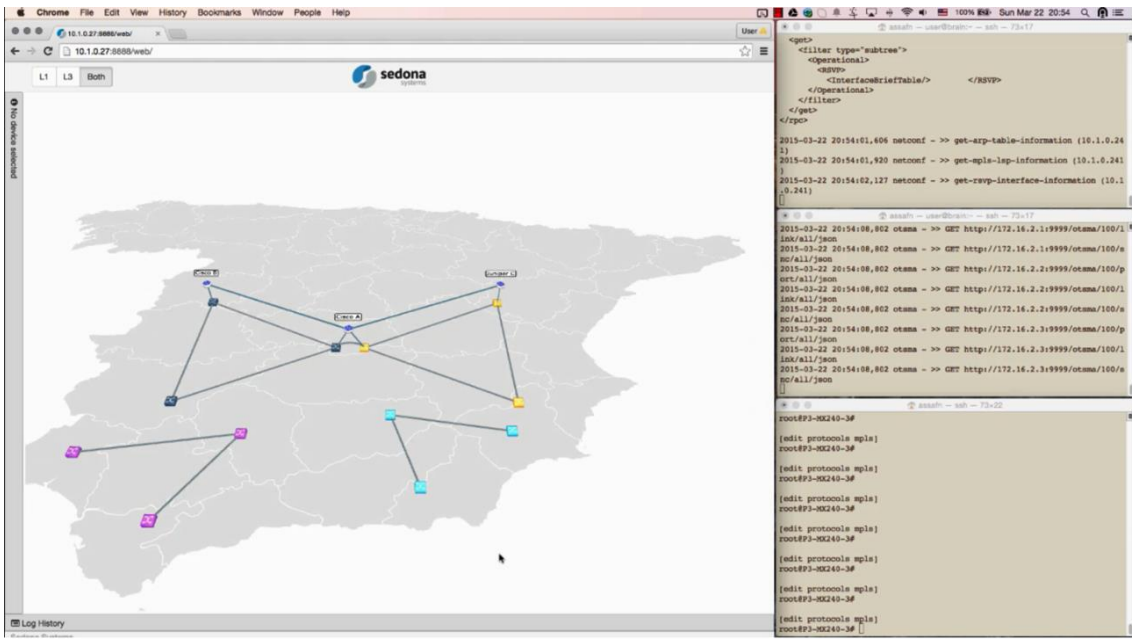

Fig. 2: Demo screenshot, including GUI of the controller, log of interactions with the routers, log of interactions with an optical controller and console of a router, used to increase LSP capacity to simulate a traffic change