
The Role of SDN in Application Centric IP and Optical Networks

Victor Lopez^{1,*}, Jose Manuel Gran¹, Rodrigo Jimenez¹,
Juan Pedro Fernandez-Palacios¹, Domenico Siracusa²,
Federico Pederzoli², Ori Gerstel³, Yona Shikhmanter³,
Jonas Mårtensson⁴, Pontus Sköldström⁴, Thomas Szyrkowiec⁵,
Mohit Chamania⁵, Achim Autenrieth⁵, Ioannis Tomkos⁶
and Dimitrios Klonidis⁶

¹*Telefónica I+D/GCTO, Madrid, Spain*

²*CREATE-NET, Trento, Italy*

³*Sedona Systems, Raanana, Israel*

⁴*ACREO, Kista, Sweden*

⁵*ADVA, Germany*

⁶*Athens Information Technology (AIT), Athens, Greece*

**Corresponding Author: victor.lopezalvarez@telefonica.com*

Received 13 September 2016; Accepted 21 December 2016;
Publication 2 January 2017

Abstract

Transport IP/optical networks are evolving in capacity and dynamicity of configuration. This evolution gives little to no attention to the specific needs of applications, beyond increasing raw capacity. The ACINO concept is based on allowing applications to explicitly specify requirements for requested services in terms of high-level (technology- and configuration-agnostic) requirements such as maximum latency or reliability. These requirements are described using intents and certain primitives which facilitate translation to technology specific configuration within the ACINO infrastructure. To support this application centric approach, SDN has a key role in this evolution. There are representative use cases where SDN gives an added value when considering not only the network but also the application layer.

Journal of Green Engineering, Vol. 6-3, 317–336.

doi: 10.13052/jge1904-4720.634

© 2016 River Publishers. All rights reserved.

Keywords: Application-Centric, IP-Optical, Multi-layer.

1 Introduction

The Internet has evolved over time into a three tier infrastructure: the top tier consists of applications driving the traffic and ultimately the requirements of the lower layers. These applications can be consumer applications such as video, audio, gaming, file-sharing, communication, social networking, consumer cloud access etc. [1], or business applications such as backup or inter-site connectivity or various datacenter-to-datacenter interactions, such as distributed search or VM migration. The application's traffic usually passes through a grooming tier, typically IP/MPLS, which aggregates multiple small flows into larger "pipes" that can be cost-effectively supported by the underlying optical transport layer. Additional grooming of IP traffic can also be performed on OTN before offloading it onto optical networks.

Traffic grooming is effective in maximizing capacity utilization and reducing management complexity. However, mapping a large number of small flows, belonging to different applications, into a small number of very large and static lightpaths means that specific application requirements, such as latency or protection constraints, are seldom guaranteed after the grooming process. While some of the requirements may be satisfied implicitly by the configuration of the infrastructure, application agnostic grooming is an obstacle to effective service fulfilment.

The ACINO (Application Centric IP/Optical Networks Orchestration) project [2] proposes a novel application-centric network concept, which differentiates the service offered to each application all the way down to the optical layer, thereby overcoming the disconnect that the grooming layer introduces between service requirements and their fulfilment in the optical layer. This allows catering to the needs of emerging medium-large applications, such as database migration in datacenters. To realize this vision, ACINO will develop an orchestrator which will expose the capability for applications to define service requirements using a set of high level primitives and intents, and then performs multi-layer (IP and optical) planning and optimization to map these onto a multi-layer service on the network infrastructure. The orchestrator also targets the re-optimization of the allocated resources, by means of an application-aware in-operation planning module.

This work presents use cases where SDN gives an added value when considering not only the network but also the application layer. The rest of the paper is structured as follows: Section 2 presents a high level explanation of the application requirements. Section 3 identifies the requirements for a

packet/optical SDN network orchestrator. Section 4 details the use cases, where a SDN orchestrator provides an added-value to applications. Finally, Section 5 concludes this paper.

2 Application Requirements

Applications requirements can vary depending on the service nature. Typical services are satisfied when there is enough bandwidth for the communication. Nevertheless, there are specific applications that may require further parameters, such as: the maximum latency, the service duration, the level of protection needed, the maximum downtime, encryption, multiple connections for the same application or even diverse routes.

The packet layer transports the traffic of multiple applications by aggregating their traffic onto optical connections. This mapping is done based on the destination address, but it is coarse nature, as the traffic is not treated according to the requirements of the application that generates it. The main reason to use the network in this way is that the granularity (and cost) of the optical connections is in the order of tens or hundreds of Gigabits per second, but the actual traffic generated by applications is typically several orders of magnitude smaller. However, two major trends are driving the change towards a different approach: on the one hand the required bandwidth for certain applications is dramatically increasing year over year and there are business applications, like Data Center Interconnect (DCI) [3], which are of a magnitude similar to that of optical connections. On the other hand, the optical layer has created mechanisms to adapt its granularity and offer a more accurate bandwidth allocation to network services [4].

Based on these premises, the idea behind ACINO is to overcome this coarse mapping by placing application-specific traffic flows directly into dedicated optical services, or, at the very least, to groom together a number of application flows with similar requirements into a specific optical service. In this manner, each application would benefit from having a transport service tailored to its specific requirements.

3 Support for Packet-Optical Networks

From a high-level perspective, this approach requires control solutions for transport networks that (a) enable applications to express their specific requirements and (b) are able to configure and reserve network resources to create a service that treats properly the application based on its requirements.

There are two main challenges when deploying services in packet-optical networks: (1) heterogeneous control planes and (2) different transport technologies inside each layer. A multi-layer SDN approach was proposed in [5] to address the same issue (Figure 1). The *orchestrator* is in charge of end-to-end connectivity provisioning, using an abstracted view of the network and also covers inter-layer aspects. Each layer has a separate *controller* that is responsible for the configuration of its own technology. Regarding the IP layer, the controller can configure several vendors, because there is interoperability. However, in the optical layer, each controller knows the vendor-specific details of its own underlying products and technologies. There are physical impairments of computation that are vendor proprietary, so each one can optimize transmission performance across the optical layer. Furthermore, the optical layer technology does not have to be the same across different optical domains. One domain can have integrated OTN switching capabilities while another domain may use WDM or even flexgrid optical switching. The only important fact to the orchestrator is that the controller offers four key services: (i) Provisioning, (ii) Topology Discovery, (iii) Monitoring and (iv) Path Computation or can easily integrate to an external application that provide such services using a simple API. Let us highlight that Figure 1 is an architectural view of the multi-layer SDN architecture. An implementation may have an IP controller and the orchestrator with the same software.

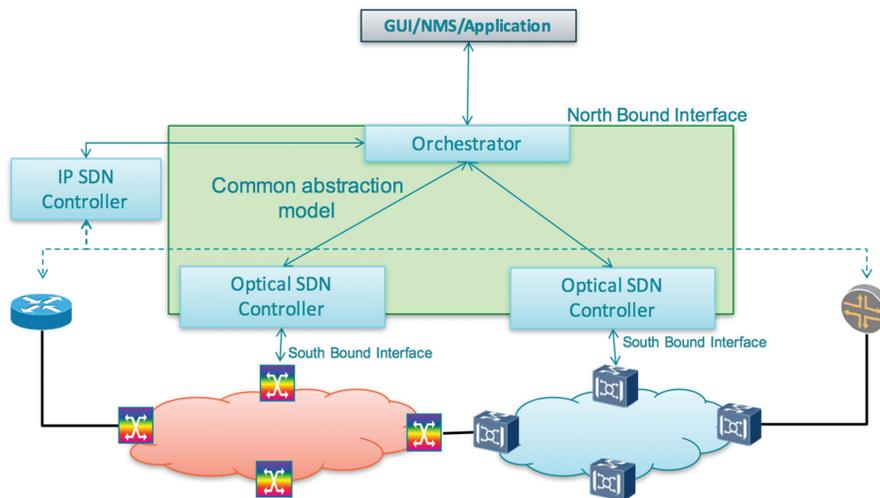


Figure 1 Multi-layer SDN architecture.

The provisioning capability enables the creation, deletion and update of connections in the network. However, to cope with any application requirement, this capability must support explicit routes, route restrictions, service resilience and traffic engineering parameters such as bandwidth and latency. Topology Discovery must export the topology information as well as the resource occupation to verify that a new service can support the application needs. The discovery of routers and optical devices is also part of this function. Monitoring capabilities are important so the multi-layer orchestrator can perform resilience actions that can not be solved locally by an underlying controller. For instance, after a failure in one domain, the orchestrator may request another connection using a different domain. Path Computation is a fundamental characteristic that allows the orchestrator to analyze candidate paths and carry out “what if” analysis. An orchestrator, with its global view of the whole network, can optimize end-to-end connections that individual controllers cannot configure.

Using these interfaces, the multi-layer orchestrator can perform the same operations as single or monovendor controllers but in a multivendor fashion. A differentiation point between the SDN and management approaches [6] is the use of standard interfaces that provides the orchestrator with a vendor agnostic view of the network resources. This would allow it to carry out multi-layer restoration operations like multi-layer re-route, which allows sharing a free IP port in each node to recover from any interface failure in the router, or Multi-Layer Shared Backup Router (MLSBR), which consists of having extra-shared backup routers to restore traffic in case of a failure of an IP router [7].

There is previous work in IP/optical demonstration like the one done in [10]. The authors uses OpenFlow to support some multi-layer use cases. The work done in [11] is more focus on commercial hardware and the authors created an optical path, configured the IP interfaces, as well as MPLS service using pseudowires. However, these tests did not use standard interfaces with the commercial hardware. Authors in [12] presented a IP/optical architecture based on IETF standard ABNO architecture, mainly relying on PCEP. There are some limitations in routers configuration that does not allows the support of multi-layer operations like MLSBR.

4 ACINO Architecture Overview

The ACINO concept is based on allowing certain applications to explicitly specify the requirements for the network services they need, in terms of high-level primitives or intents; in this document, we use the following definitions:

- Primitive: a parameter that expresses a service feature, e.g. network object or constraint;
- Intent: a combination of desired primitives that corresponds to the desired characteristics of a useful service.

In other words, in line with existing literature on the subject, we refer to “primitives” when describing one or more features of the underlying network(s) that the ACINO orchestrator is able to handle, and we call “intents” to the combination of primitives that a client application selects to describe the desired characteristics of a service.

At a high level, the ACINO architecture is based on a logically centralized **Orchestrator**, outlined in Figure 2, whose primary function is translating the application-level service requirements into appropriate configuration requests for both the underlying IP/MPLS and Optical network layers. The orchestrator interfaces with the outside world through two main interfaces. Firstly, via a **North-Bound Interface** (NBI) towards applications, which gives them the possibility to request network services with specific requirements, such as latency, reliability and capacity. Special applications, such as a Network Management System (NMS), may even interface with the online planning functionalities exposed by the orchestrator. These “application intents” are then translated into a multi-layer service configuration utilizing both IP/MPLS and optical resources, which is pushed down to the underlying network layers through the use of a **South-Bound Interface** (SBI), which connects the orchestrator to the control planes of both the optical and the IP networks, and is used to both push configurations and pull network state or relevant alarm states from the networks being controlled.

Therefore, the ACINO orchestrator is a platform that performs the following functions: (i) it gathers information on the underlying network layers using the SBI and stores it into a multi-layer network model; (ii) it receives

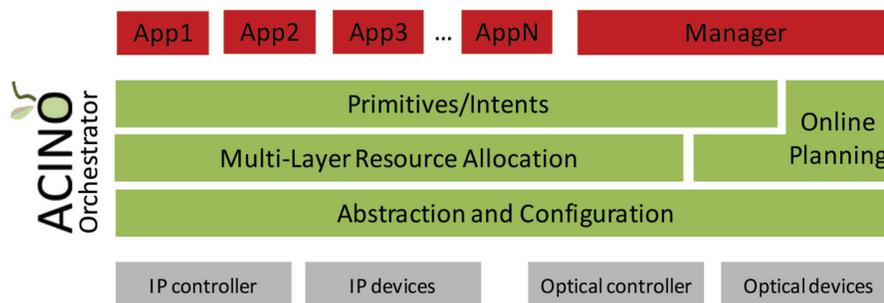


Figure 2 Simplified schema of the ACINO orchestrator.

service requests from applications through the NBI in the form of “intents”, maps these intents into installable configuration for the underlying networks, optimizing it with respect to both the explicit service requirements (which must be satisfied, if at all possible) and existing configuration (re-using pre-established lightpaths and Traffic Engineering tunnels whenever possible), and pushes such configuration down to the underlying network control planes via the SBI, and (iii) it receives planning requests from one or more Network Management System class application through the NBI, computes the effects of the proposed changes and replies back to the requestor with the outcomes of the computation; optionally, it may also push the resulting configuration to the underlying network layer via the SBI. Similarly, the ACINO orchestrator will receive alarms from the network elements and it will react to fulfill the applications needs, based on the available resources and re-configure the network. A demonstration of the restoration capabilities of the ACINO orchestrator was done in [13].

5 Use Cases

The ACINO consortium has selected some applications and relevant network operations to illustrate, with use cases, the value proposition of SDN in application centric IP/optical networks. Note that ACINO is currently envisaged and designed for core and metro networks typically spanning large geographical areas. Therefore, individual consumer end-to-end applications typically running on hosts are not expected to benefit from the ACINO approach, due to it covering a limited part of the network path, and requiring additional overhead to set-up or select a customized service (although the application-centric paradigm could be extended to better cater to these applications). The main beneficiaries of the ACINO approach are business applications with very specific (and known) needs (including security, latency, availability constraints), such as e.g. distributed banking systems, Content Delivery Networks (CDNs), applications spanning multiple data centers or others. With respect to bandwidth granularity, ACINO supports bandwidth requirements as one of the many characteristics of a desired service, but it caters to everything from very small to very large demands.

5.1 Application-based DataCenter Interconnection

An initial case study for the ACINO approach is that of optimizing multiple large network-facing applications requiring a services with specific characteristics. For example, one business application may need to migrate VMs

according to a “follow the sun” approach, which entails regular, schedulable, large-sized, short lived network flows for which latency is not paramount. Another application, consisting of a distributed, synchronously updated DB, may require a constant low bit rate connection with minimal latency. The owner of a third application may suddenly decide to move all related VMs to a different DC, for example because the current ones lacks infrastructure to support future expansion or to maintain availability during extraordinary maintenance to the original DC. This is a one-time, possibly schedulable event involving the bulk transfer of a large amount of inter-dependent VMs, and would therefore need a trade-off between bandwidth and duration, with a likely emphasis on the former.

Assuming that these applications are housed in the same DC and the first two need connectivity to the same external DC, while the third one to a different DC, as shown in Figure 3, currently the first two would be mapped to the same service (with shared characteristics), potentially sharing an IP TE link with the service used for the third application.

Using the ACINO approach, these applications (or some human operating on their behalf) could explicitly specify their requirements to the ACINO orchestrator to map them in appropriate services. Let’s assume that the two DCs are already connected using a L3 VPN with no special requirements save some guaranteed bandwidth, as shown in Figure 3.

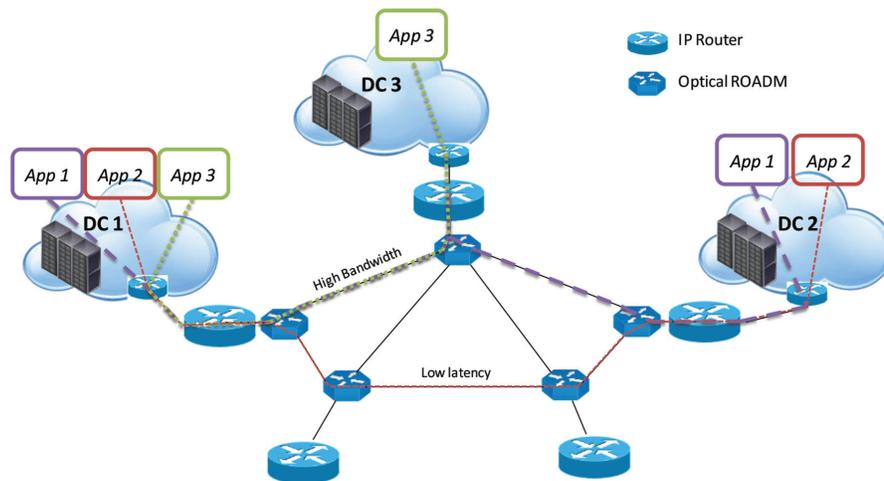


Figure 3 Application centric Intent-based IP optical orchestration.

The first application just requires a large amount of bandwidth periodically. The orchestrator could decide to periodically set up another, dedicated, WAN link to satisfy the application, or it could reserve extra bandwidth on the existing VPN link (e.g. if the underlying optical connections are overprovisioned with respect to configured tunnels), or even do nothing at all if the baseline service is deemed sufficient.

The second application requires a constant connection with low latency. If the baseline service is not already configured as an IP adjacency served by one (or more) optical connection(s) on the physically shortest path, the orchestrator would set up such a service and instruct the border router to direct traffic from that application over it.

Finally, the third application requires a large amount of bandwidth on a potentially short notice. Unless the baseline service is largely overprovisioned, the orchestrator might need to temporarily assign more capacity to it, by leveraging available bandwidth in the optical layer or possibly instantiating new temporary optical connections. Furthermore, since, overall, two datacenters need to be connected to the original one, the orchestrator may decide to share a single optical adjacency for both connections up to an IP router near one of the datacenters, and simply re-send the traffic destined for the other DC to the optical layer to be carried on a second optical connection.

5.2 Enabling Dynamic 5G Services

Dense deployments of small cells, often referred to as ultra-dense networks (UDN), will be a major building block for increasing capacity in 5G radio access. The scenario considered here is a UDN deployed in a location where a large crowd of people is gathered in a relatively small area for limited time periods, where all the small cells are switched off most of the time. One example is a stadium where the UDN is active for the duration of a live event, such as a football match or a concert. A stadium UDN deployment would benefit from the ACINO approach through the ability to dynamically request and reserve backhaul transport capacity for when the UDN is active.

A key vision for 5G is that the network should support a high variety of applications with very diverse requirements. For this reason, the 5G service deployment concept of network slicing has emerged, where one physical infrastructure is shared between all these applications by introducing multiple virtual and logically separated instances spanning all network domains, including the transport network. The ACINO approach would enable extending this differentiated 5G application treatment to the optical layer.

Figure 4 illustrates how the ACINO transport network could support the stadium scenario. A mobile edge cloud is located at or close to the stadium, running selected distributed 5G user plane (UP) or control plane (CP) network functions depending on network slice. The UP may include a traffic separation function ensuring that uplink traffic belonging to different slices is separated, while typical CP functions are session authentication and policy control. From this edge cloud, multiple transport network connections are set up to support the different network slices and applications.

Figure 4 depicts three application examples with different requirements. One application, expected to dominate the traffic volume in the stadium scenario, is upload/download of videos and photos from members in the audience to/from a central cloud (e.g. YouTube). This application would have high requirements on bandwidth but less focus on latency and reliability. Another application is the remote control of cameras filming the event and transmission of their video streams to a central production studio. This application would have requirements on low latency and high reliability. The third application is connecting a number of different environmental sensors at the stadium to a central location with low requirements on bandwidth and no strict latency requirements.

Note that in Figure 4, even though all three applications enter and exit the ACINO transport network at the same nodes, they take different paths through the network. Of course, all traffic does not have to pass through the same exit node; more latency sensitive applications may e.g. be directed to a more local datacenter, and some applications may even run on servers in the edge cloud located at the stadium.

5.3 Application-Specific Protection Strategies

Optical layer restoration and multi-layer restoration have been researched extensively and they can account for very substantial savings in the total number of required router interfaces and transponders, on the order of the 40–60 percent in the core, as shown in Figure 5 [9]. In both cases, it is assumed that some of the responsibility for restoring from failures is moved to the optical layer, since it is much more cost effective to build in spare capacity in the optical layer than to do so in the IP layer. The former approach assumes that the optical layer alone is responsible for restoring from a failure – and therefore the selection of the restoration path is insensitive to the needs of the IP layer. The latter approach does take IP layer needs into account, but for the aggregate traffic that traverses the failed IP links.

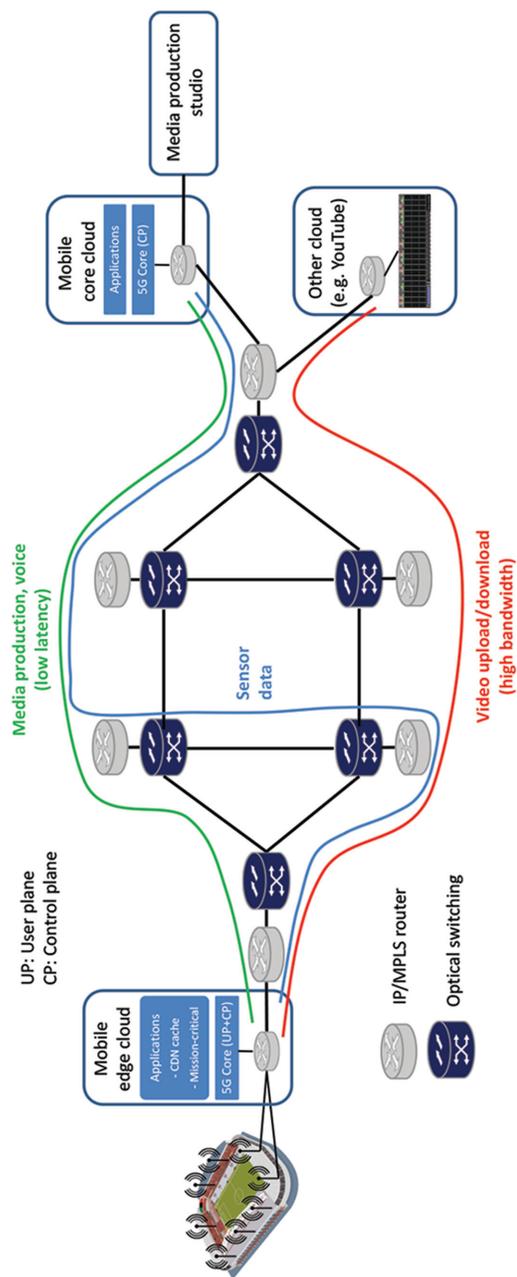


Figure 4 Illustration of the dynamic 5G services use case. In this scenario, the ACINO network provides dynamic transport capacity to a stadium and three different 5G applications with different requirements are depicted.

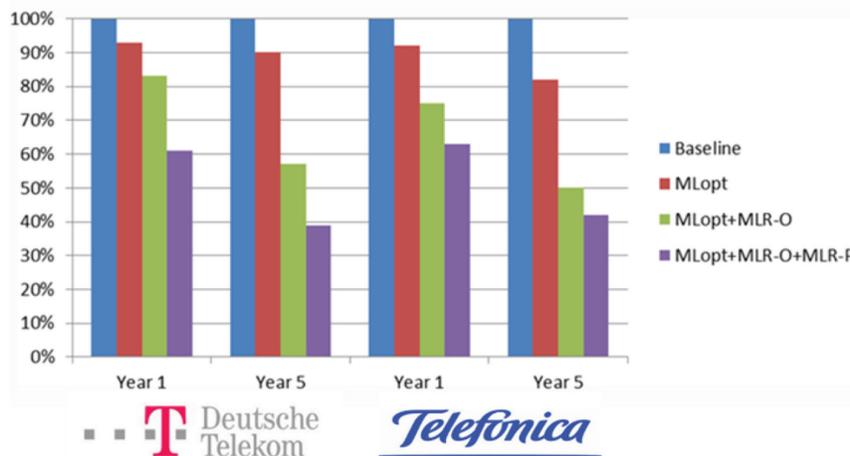


Figure 5 Savings due to multi-layer strategies.

An example for the behaviour of a latency-sensitive traffic can be found next. In this example, we assume all optical links have the same length and that a service can tolerate a latency of 4 optical links. Figure 6a shows a service routed over this network in green (on the left), and its routing over the network after the optical recovery from a failure, assuming it takes the same IP layer path. This is acceptable for non-latency-sensitive traffic, however if the max latency is 4 hops, then the IP layer should route the service over a different IP path – as shown in Figure 6b. If the protection would be only done at the optical layer the service would be rerouted as shown in Figure 6c. However, there is an increment on the delay as the service has to cross the whole ring. Therefore, the ACINO orchestrator will perform an application-specific restoration to provide each application with the required resilience mechanism, as depicted in Figure 6d.

5.4 Secure Transmission as a Service

The frequency of cyber-attacks on critical network infrastructure and public/private entities connected to the Internet has been growing significantly and has translated into significant financial costs for end-customers or businesses. As a result, more and more businesses are investing in improving their security infrastructure. Moreover, server-to-server communication, especially data-center interconnects (DCI) are increasingly becoming an important service offering for network operators.

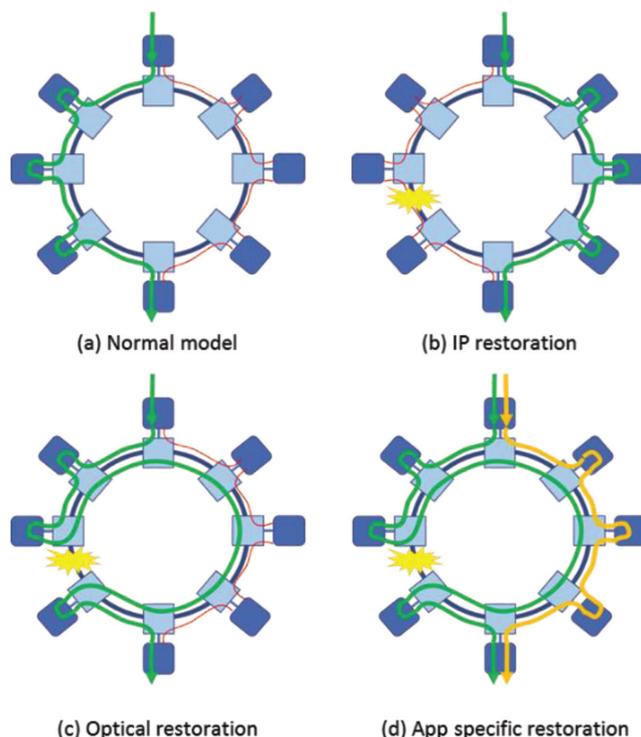


Figure 6 Example for restoration under application latency constraints.

One of the most common trends for securing the communication has been to push for end-to-end encryption (IPSec, HTTPS). End-to-end encryption is flexible and independent of the underlying network infrastructure, making it relatively easy to deploy. However, the flexibility comes at the cost of increased processing requirements at both server and client endpoints, increased latency and reduced throughput from the network. Shifting the responsibility to the network reduces the processing complexity at the endpoints (servers) and allows network operators to optimize the \$/bit cost for encrypting traffic between two remote sites. Different mechanisms such as IPSec, IEEE MacSec, and custom all-optical encryption differ on the cost of deployment, availability, latency and throughput.

The distinctive features of these technologies provide an opportunity for the network operators to differentiate their service portfolio to meet the needs of specific applications. For example, for dynamic low volume interconnects between two or more sites in a meshed configuration, secure connectivity

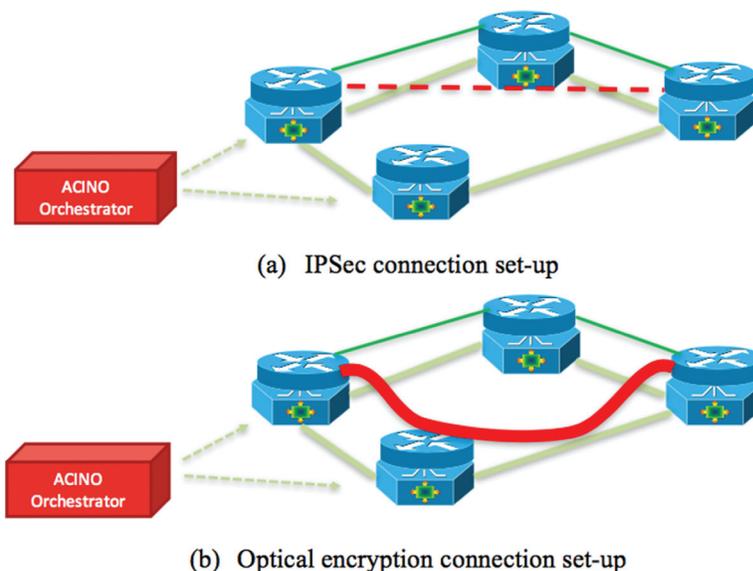


Figure 7 Secure transmission as a service examples.

can be established by setting up IPsec tunnels between the sites (Figure 7a), while applications such as interconnects between financial institutions, that are extremely sensitive to latency, or large dataset transfers, that are extremely sensitive to throughput, can employ all-optical encryption (Figure 7b) that requires custom infrastructure (and is therefore a limited resource).

To demonstrate this use case, the REST NBI of a controller can be extended to support the encryption, as any other constraint needed for a service (like latency or bandwidth). The authors in [14], extended ONO's NBI with an encryption capabilities to indicate that the encryption processing has to take place, otherwise the intent is processed like an unencrypted request. Based on this extension, this use case can be demonstrated.

5.5 Dynamic Virtual CDN Deployment

Content delivery networks (CDN) are used to deliver content from servers located in datacenters (DCs) to end users. In this case, we will consider video distribution as a reference application for the CDNs. Videos are delivered by using the IP layer. Therefore, the choices for an operator's CDN location are limited to the IP core. Possible deployment sites include regional DCs at the access routers (AR), high density DCs at the transit routers (TR) or

even national DCs at the interconnection level (IX). Deployment of video applications only on National CDN has the advantage of high utilization of the video servers, since users access the same datacenters, thus statistically using the resources more often. However, each connection or video between the user and one of those national video servers would pass the whole national network. As the data flows are unicast for video on demand or time shifted services, the (redundant) overhead in the network would be massive. On the other hand, the video servers deployed at the AR locations minimize the network overhead, but increase the video platform's costs by increasing the number of locations and redundant copies of the same content. Moreover, if the number of customers using these video servers is small, the dimensioning and caching is not efficient.

As the traffic nature is dynamic (events in stadiums, popular contents in regions, unexpected high penetration . . .), the use of a virtual CDN infrastructure makes sense. A virtual CDN infrastructure has the same capabilities as current CDN deployments, but it runs in standard x86 servers. This way the vCDN provider can dynamically activate VMs in locations that require more video servers (Figure 8). Once the video server is activated, the contents can be transferred from a national/regional datacenter to sync the most popular content. This requires an inter-datacenter transfer, which in turn requires

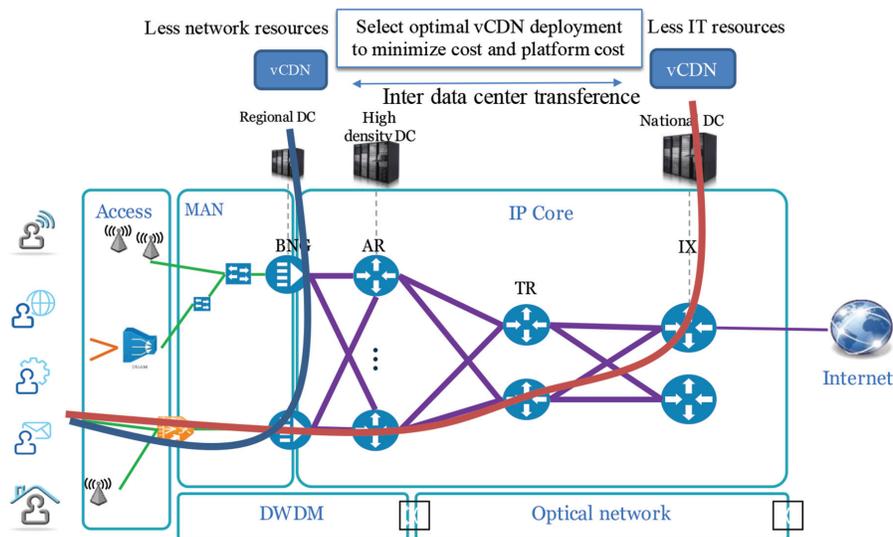


Figure 8 Virtual CDN scenario.

network resources for a one-time sync between the fixed and the virtual CDN site. This approach can reuse the investment in datacenters that operators are doing for virtualizing services or even Virtual Network Functions (VNFs).

In the vCDN scenario, there are two main traffic flows: inter data center traffic and data center to end user traffic. To perform this use case on a real network, IP+optical coordination is critical to optimize the network resources. There are some hours, where the network is under utilized. The orchestrator must use this time to set-up direct connections between the data centers to carry out the content synchronization. Similarly, when the inter-data center traffic is low these resources must be used for the data center to end customer traffic reusing the IP cards and transponders.

5.6 Application-Centric In-Operation Network Planning

Current transport networks are statically configured and managed, because they experience rather limited traffic dynamicity. This leads to long planning cycles to upgrade the network and prepare it for the next planning period. The planning procedure aims to ensure that the network can support the forecast traffic and deal with failure scenarios, thus adding extra capacity and increasing network expenditures. Another drawback of current static procedure is that the results from network capacity planning are manually deployed in the network, which limits the network agility. The main reason for this is that current provisioning systems are not always deployed as the number of configurations in IP/optical backbone networks is relatively small. The authors of [9] proposed the term “in-operation network planning”. The main idea of in-operation network planning is to take advantage of novel network reconfiguration capabilities and new network management architectures to perform in-operation planning, aiming at reducing network CAPEX by minimising the over-provisioning required in today’s static network environments.

The overall framework architecture of a statically configured network and a network that supports in-operation planning is shown in Figure 9. The key differentiation for the adoption of the in-operation planning approach is in the undertaken operations of the network provisioning system, which should now facilitate an in-operation (i.e. dynamic and real time) planning tool that interacts directly with the data and control planes. The adoption of a SDN framework with north and south bound interfaces supports the realization of this concept in a multi-service and multi-vendor environment.

Within ACINO, the in-operation planning concept is extended and includes application awareness. This denotes that overall planning of the resources is

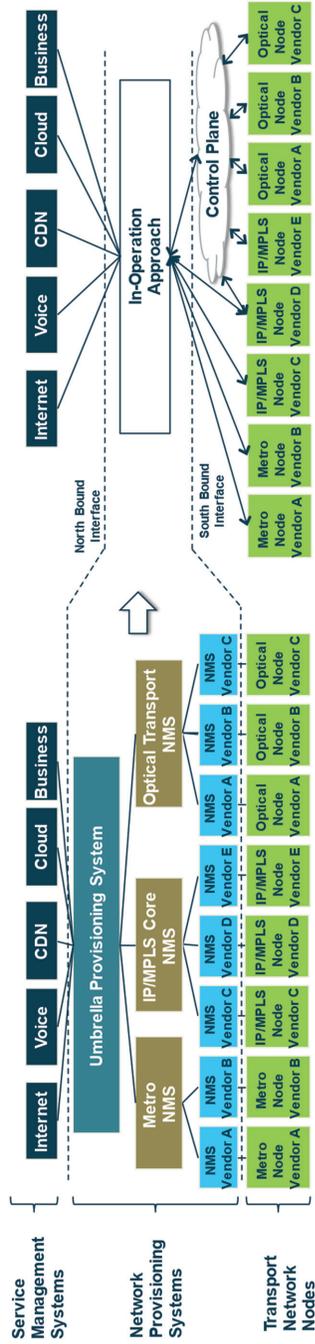


Figure 9 Evolution towards in-operations planning.

not performed solely for the optical infrastructure considering the aggregated data from the upper layer, but on a ‘per application demand’ basis considering both the IP and Optical layer resources. With the ACINO dynamic multi-layer approach, the incoming requests are classified in a sense that they generate different service requirements to the network which translate to different use of the available resources. These requirements are evaluated in real-time providing the optimal routing through IP and optical domain or establishing new IP or/and optical lightpaths if this is required.

6 Conclusions and Future Work

The ACINO project is built on top of the idea that, today, application traffic is aggregated in the network infrastructure. Therefore, it receives a common treatment in different network segments even for critical parameters like bandwidth, latency, restoration or resilience, whose actual requirements may well differ between different applications. By applying multi-layer resource optimization algorithms with an application-centric approach, the applications can receive a network service tailored to their requirements, while the network resources can be appropriately allocated. Indeed, the packet layer provides a transport to multiple applications by aggregating their traffic onto optical connections. However, this mapping is done on the basis of the destination address, which is naturally coarse, so the traffic is not treated according to the requirements of the application that generates it. The value of a more direct mapping of application flows onto the network services can be significant, because it allows the specific requirements of the application layer to be directly mapped onto the packet and optical layers.

This paper details the use cases where SDN can play an important role for Application Centric IP and Optical Networks. It presents an analysis of the application requirements as well as the key elements to support the network operations for packet-optical scenarios. As future work, the authors will analyze each case study to validate the improvement of the application centric approach in packet-optical networks. Moreover, the ACINO project will assess CAPEX and OPEX parameters like Energy Efficiency. Energy reduction is expected due to cross-layer application-aware offline and in-operation network planning as well as due to the application of effective online resource allocation algorithms to efficiently use the network resources, based on application traffic prediction and forecasting.

Acknowledgment

The research leading to these results has received funding from the European Commission within the H2020 Research and Innovation program, ACINO project, Grant Number 645127, www.acino.eu

References

- [1] Sandvine Global Internet Phenomena Report, 2H (2013). *Sandvine Global Internet Phenomena Report, 2H*. Available at: <https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/2h-2013-global-internet-phenomena-report.pdf>
- [2] ACINO H2020 EU Project. Available at: http://www.acino.eu/wp-content/uploads/2015/02/ACINO_Public-Project-Presentation.pdf
- [3] Mahimkar et al. (2011). “Bandwidth on demand for inter-data center communication,” in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks, HotNets-X*, Cambridge, MA.
- [4] Gerstel, O. (2010). “Flexible use of spectrum and photonic grooming,” in *Proceedings of the Photonics in Switching 2010*, Monterey, CA.
- [5] Gerstel, O., and López, V. (2015). “The need for SDN in orchestration of IP over optical multivendor networks,” in *Proceedings of the European Conference on Optical Communication* (Rome: IEEE).
- [6] Martínez, A., Yannuzzi, M., Lopez, V., López, D., Ramírez, W., Masip-Bruin, X., et al. (2014). Network management challenges and trends in multi-layer and multi-vendor settings for carrier-grade networks. *IEEE Commun. Surv. Tutor.* 16, 2207–2230.
- [7] Mayoral, A., López, V., Gerstel, O., Palkopoulou, E., Fernández-Palacios, J. P., and González de Dios, Ó. (2014). Minimizing resource protection in IP over WDM networks: multi-layer shared backup router. *Proc. Opt. Fiber Conf.* 7:M3B.1.
- [8] Gerstel, O., Filsfils, C., Telkamp, T., Gunkel, M., Horneffer, M., Lopez, V., et al. (2014). Multi-layer capacity planning for IP-optical networks. *IEEE Commun. Mag.* 52, 44–51.
- [9] Velasco, L., King, D., Gerstel, O., Casellas, R., Castro, A., and Lopez, V. (2014). In-operation network planning. *IEEE Commun. Mag.* 52, 52–60.
- [10] Das, S. (2012). *pac.c: Unified Control Architecture for Packet and Circuit Network Convergence*. Ph.D. thesis, Stanford University, Stanford, CA.
- [11] Muñoz, F., Muñoz, R., Rodríguez, J., López, V., González de Dios, O., and Fernández-Palacios, J. P. (2013). “End-to-end service provisioning

across MPLS and IP/WDM domains,” in *Proceedings of the International Workshop on Network Management Innovations Co-located with the 4th IEEE Technical Cosponsored International Conference on Smart Communications in Network Technologies* (Rome: IEEE).

- [12] Aguado, A., López, V., Marhuenda, J., González de Dios, Ó., and Fernández-Palacios, J. P. (2015). ABNO: a feasible SDN approach for multi-vendor IP and optical networks. *J. Opt. Commun. Netw.* 7, A356–A362.
- [13] Santuari, M., Szyrkowiec, T., Chamania, M., Doriguzzi-Corin, R., López, V., and Siracusa, D. (2016). “Policy-based restoration in IP/Optical Transport Networks,” in *Proceedings of the IEEE NetSoft Conference and Workshops* (Rome: IEEE).
- [14] Szyrkowiec, T., Santuari, M., Chamania, M., Siracusa, D., Autenrieth, A., and López, V. (2016). “First demonstration of an automatic multilayer intent-based encryption assignment by an open source orchestrator,” in *Proceedings of the Post-Deadline Paper in European Conference on Optical Communication*, Frankfurt.