ACINO

Application-Centric IP/optical Network Orchestration

ACINO / MS31

ACINO concept paper

Dissemination level	PU
Version	1
Due date	31/01/2016
Version date	29/03/2016

This project is funded by the European Union

1 Introduction

Applications and verticals that support specific business processes, with their peculiar requirements and targets, will ultimately drive the evolution of telecommunications towards 5G and beyond. In the transport networking area, relevant business applications could be, for example, virtual machine migration, data backup (e.g. sensitive financial or social network data), distributed content delivery (e.g. video-on-demand, TV broadcasting), etc. In the last few years, they have evolved from simple requirements, that can be easily and cheaply met with a best-effort network and a reliable end-to-end transport protocol, to stringent requirements in terms of latency and jitter (e.g. financial transactions), bandwidth (e.g. video production over IP), reliability, security, etc.

Such applications operate on top of a three-layer structure: they generate traffic that is groomed at the IP/MPLS and/or OTN layers and finally transported at the optical layer. Despite these diverse requirements, the application's traffic usually passes through a grooming layer, typically IP/MPLS, which aggregates multiple small flows into larger ones that can be cost-effectively supported by the bottom optical layer (sometimes there is also OTN, another grooming layer between the IP and photonic layers). The existence of these grooming layers results in an indirect relationship between the application layer and the transport layer, which implies an inaccurate mapping of application needs to the transport layer, since the requirements imposed by the grooming layer to the optical layer are the result of an aggregation of diverse applications with different needs. While offering the best possible service characteristics at the optical transport layer to each application is theoretically achievable (e.g. by using a fully meshed optical bypass virtual topology), it would be prohibitively expensive and energy-hungry. Therefore, a smarter approach is needed.

The ACINO (Application Centric IP/optical Network Orchestration) project proposes an application-centric approach, where the traffic of each application receives a tailored service at each layer of the transport network (potentially, all the way down to the optical layer), thereby overcoming the gap that the grooming layer introduces between application service requirements and their fulfillment in the lowers layer of the stack. This way the network becomes more efficient, since the grooming process is minimized and the optical resources are used efficiently to meet but not needlessly exceed the requirements of the applications. The applications, in turn, could exploit – and even be engineered to take advantage of – the on-demand and personalized reservation of network connectivity, resources, reliability, security, etc. As a result, the needs of emerging medium-large bandwidth-consuming applications will be catered to, while small flows that do not match the granularity of an optical service can still be groomed together with other flows having similar requirements (i.e., classes). The ACINO concept is outlined in Figure 1.

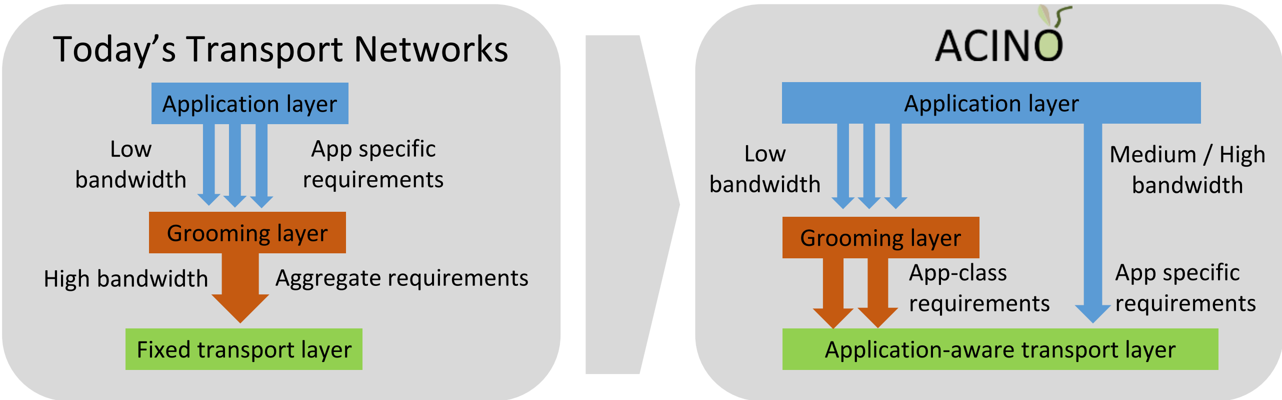

Figure 1. ACINO concept.

ACINO focuses on the control aspects of networking, while trying to achieve its outcomes by leveraging the state of the art of transport technologies (both IP and optical). For this reason, the project exploits the Software-Defined Networking (SDN) approach in the context of multi-layer dynamic network control. Indeed, with its strong focus on the programmability aspects, SDN can produce evident benefits on the network operations and provide a powerful interface towards the applications, in line with the needs of the project. Such a joint control of different sets of network technologies spanning multiple layers has been referred by the project to as *network orchestration*.

In order to bring the envisioned concept to reality, the ***objective of ACINO*** is to develop and experimentally demonstrate an *open-source IP/optical network orchestrator* that includes online planning and computation software modules exploiting advanced application-centric methods and algorithms, and exposes specific primitives towards the applications to easily map their needs to the service they will receive.

This concept paper is organized as follows. Section 2 summarizes and briefly describes the applications that have been identified as most relevant, together with their specific requirements. Section 3 illustrates the case studies that have been considered as the most promising to demonstrate the innovation produced by the ACINO project. Section 4 gives a high-level description of the main components of the ACINO orchestrator, while Section 5 concludes the concept paper.

2 ACINO applications and requirements

As already explained in the previous Section, the traffic of multiple applications is aggregated onto optical connections. Such mapping is coarse by nature, as the traffic is not treated according to the requirements of the application that generates it. The main reason behind this inefficient usage of resources is that the granularity 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 some orders of magnitude smaller. However, new trends have been emerging in the very last years. First of all, the required bandwidth for the applications is dramatically increasing and there are business applications, like datacenter to datacenter, for which bandwidth requirements are not small anymore with respect to the granularity of the optical connections.

Secondly, the optical layer has adopted mechanisms to adapt its granularity and offer to the services a more accurate bandwidth allocation.

According to the current trends, three groups of applications could mostly benefit from the ACINO concept:

- **5G applications** (e.g. enhanced mobile broadband, massive machine-type communications, fronthaul services to enable Cloud RAN deployments): depending on specific needs, they are usually characterized by a combination of high throughput, low latency and high reliability requirements.
- **Video applications** (e.g. IPTV, Video on Demand, OTT video): they can have diverse bandwidth and latency requirements, and may require encryption and multicast services.
- **Cloud applications** (e.g. virtual machine migration or replication, database synchronization): they can have different latency requirements (for example, applications requiring synchronization are very latency sensitive) and are usually very demanding in terms of peak capacity.

In order to satisfy the outlined requirements, the network resources should be configured and adapted accordingly, as it will be explained in Section 4. In the next Section, we describe some case studies, given by a combination of applications that receive a tailored service by resorting to one or more multi-layer network operations.

3 ACINO case studies

Six case studies that can help demonstrating the innovations produced by the ACINO project have been identified. Such case studies show how the adoption of the ACINO concept can help catering the needs of the applications in multiple scenarios by enforcing the most proper configuration both at the IP/MPLS and optical layers and by executing operations such as provisioning, in-operation planning, survivability, encryption.

The case studies are described in the following:

- **Application-based data center interconnection**: this first case study has been designed in order to provide a basic demonstration of the value of the ACINO concept. Three applications that need to move VMs or synchronize data across different data centers are considered. These applications have different requirements in terms of execution time, frequency and required bandwidth. The application of the ACINO concept will allow to cater the needs of the applications without over-provisioning the network.
- **Enabling dynamic 5G services**: this case study aims at demonstrating that the ACINO concept can be applied in a 5G scenario, where an Ultra-Dense Network (UDN) is deployed in a location where a large crowd of people is gathered in a relatively small area for limited time periods, and where all the small cells are switched off the rest of the time. The application from ACINO's perspective is then to dynamically provide backhaul capacity for the UDN when it is active, either as a single application, or as multiple applications by offering differentiated services for different traffic flows (e.g. a UDN

deployed for the duration of a live event in a stadium, the backhaul transport capacity is dynamically requested and reserved).

- **Application-specific protection strategies**: this case study proves the benefits of applying the ACINO concept in a resiliency scenario, considering the difference of providing IP or optical restoration for latency-sensitive applications.
- **Secure Transmission as a Service**: a relevant field to explore is that of network infrastructure security, which has grown significantly in the recent years and has translated into significant financial costs for end-customers and companies. With ACINO, different applications may benefit from secure transmission services applied at different layers, implying different costs, availability, latency and throughput.
- **Dynamic Virtual CDN deployment**: this case study demonstrates the convenience, for certain applications (e.g. according to the popularity of the contents), of dynamically placing and activating virtual CDN servers in order to deliver the video contents to the end-customers. This approach will also allow the re-use of investments in data centers that operators are doing to virtualize their services.
- **Application-centric in-operation network planning**: in-operation planning has been already introduced in other projects, but with ACINO this concept is extended and includes application awareness. This case study is devoted to demonstrating the importance of considering applications' requirements when performing such operations.

In the next Section we will focus on the definition of the ACINO network orchestrator, which is in charge of carrying out the network operations required by each case study discussed so far.

4 ACINO solution: the network orchestrator

The ACINO project aims at defining the overall architecture and the requirements of the multi-layer (IP and optical) network orchestrator, as well as its interfaces and supporting network controllers. Moreover, a set of requirements for the controllers of the IP layer and of the optical layer that will enable multi-layer control needs to be defined.

A high-level overview of the main elements of the ACINO orchestrator platform is shown in Figure 2. The orchestrator interfaces with the outside world through two main interfaces. A North-Bound Interface (NBI) towards applications gives them the possibility to request network services with specific requirements such as latency, reliability, capacity etc., in the form of "intents". Special applications, such as a Network Management System (NMS), may even interface with online planning functionalities exposed by the orchestrator. Application intents provided over the NBI 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 South-Bound Interfaces (SBI) that are exposed by the control planes of the optical and IP networks. The SBIs are used to push configurations to the network and to pull the network state or relevant alarm states from the networks being controlled.

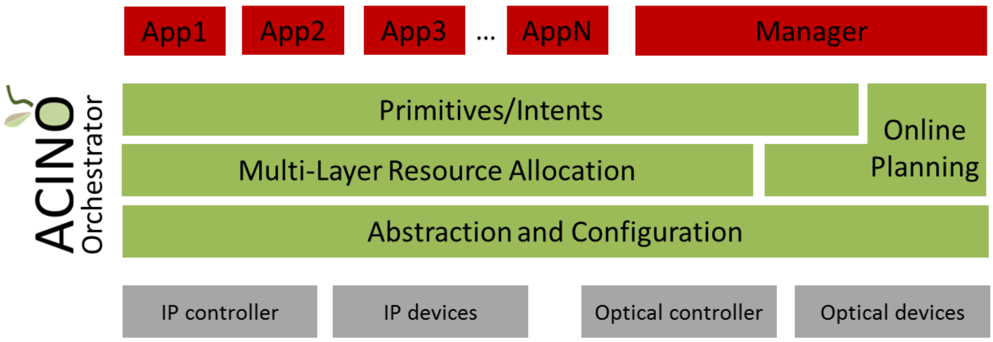

Figure 2. Simplified schema of the ACINO orchestrator.

More specifically, the orchestrator platform performs the following operations: (A) it gathers information on the underlying network layers using the SBI and stores it into a multi-layer network model; (B) it receives intents from applications through the NBI, maps them into an installable configuration for the underlying networks, and pushes such configuration down to the underlying network control planes via the SBI; (C) it receives planning requests from one or more Network Management System class applications 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.

In the following, we will better explain how heterogeneous applications can request network services with specific requirements in the form of intents and how the orchestrator can efficiently translate such requirements in network configurations to be pushed to the optical and IP networks.

4.1 Intent-based North-Bound Interface

A common definition of the key parameters that allow applications to explicitly specify requirements for the network services, without delving into the details of the underlying technologies, is needed. This way, the implementation of the application request is up to the orchestrator, which has full freedom to decide how to provide the requested service.

Based on what can be realistically provided by networking hardware, the ACINO project proposes the definition of an application-centric intent-based NBI. This interface allows the application to discover a number of Primitives and a Grammar describing their valid combinations. The application can then choose the best combination of primitives (i.e., intent) to best express its service requirements. Application-specific behaviors can thus be described by multiple primitives, such as Bandwidth, Latency, Security, Cost, and Availability. All these primitives are agnostic of the underlying network technology, not restricting the NBI to be used only on top of a specific network technology.

Once an intent has entered the orchestrator, the Primitives/Intent module (see Figure 2) is responsible for compiling it into possible solutions: this task is performed by the module through the validation and resolution of the different primitives in the intent. Additional operation can also be performed. For example, if the intent contains priority primitives, the solutions can be ranked according to them; if no solution satisfying all the constraints of the incoming intent is found, then a negotiation with the application may begin, offering one or more solutions that partially satisfy the request and letting it decide which, if any, best suits its needs.

4.2 Strategies for Application-Centric Multi-Layer Resource Allocation and Planning

Once one or more possible solutions are derived, they are fed to the Multi-Layer Resource Allocation module or to the Online Planning module (see Figure 2). Such modules map them to the network topology and calculate their costs and performance values (e.g. in terms of latency); they ensure that the network resources needed by the applications are available and the network is run efficiently. To that end, the modules perform algorithmic network optimization that relies on multi-layer strategies.

In addition to application service requests, the optimization process must take into account a particular internal cost function (such as minimum wavelength or power usage), the features of the available equipment and the current network configuration. The orchestrator views the underlying network as two-layered: optical- and IP-based. Both layers have their own requirements and unique features that, if exploited, allow for a more optimized network.

Concerning the Multi-Layer Resource Allocation module, the general strategy for multi-layer resource allocation is to reuse existing routes in the network, whenever possible. Clearly, whenever existing routes do not satisfy the requirements of an incoming service request the optimization modules will resort to setting up new routes. The two network layers and the two main strategies (reuse a route vs. set up new route) in each layer thus define four basic actions that the provisioning and planning modules can take.

Moreover, the Online Planning module is in charge of re-optimizing the network. Periodically, if the state of the network is such that it could benefit from moving some traffic or lightpaths, the module calculates which to move where and how. A further function of the module is to perform what-if calculations, useful to the network operator. Specifically, it calculates the necessary network resources needed for a future set of requests, and the effects of network upgrades or failures.

Finally, once decisions are taken by the Multi-Layer Resource Allocation and Online Planning modules, the application-centric orchestrator must be able to configure and correlate multiple underlying network layers through the South-Bound Interface. To do so, it exploits a network model that is able to capture the needed details of each and every network layer being controlled.

5 Conclusion

The ACINO project aims at addressing the mismatch between the service provided by the transport layer and the needs of verticals/applications, by creating multi-layer services tailored to the requirements of the applications. This way, verticals and applications owners will definitely benefit from a transport ecosystem that is capable of “talking their language”, service providers will be able to offer more cost effective services and network operators will be able to transfer, protect and secure large amounts of data generated by applications in a more effective way. Moreover, the ACINO project will provide strong motivation for increasing the capabilities of the network, giving system vendors and subsystems/component vendors an advantage in building an alternative food chain geared towards more flexible solutions.

With respect to the applicability of the ACINO concept, this is not limited to backbone transport networks: in fact, it can ensure that specific 5G requirements can be satisfied end-to-end, on a per-application basis. From the operational perspective, this can be done either by coordinating the ACINO network orchestrator with the controller and resource manager of the mobile network, or by adopting (and adapting) it for the transport segments of the 5G network (i.e., fronthaul and backhaul).

To summarize, the expected impact of ACINO can be decomposed into the following points:

- Enable transport network services directly satisfying the application's needs;
- Prove the value of dynamic and elastic optical technologies;
- Tackle the lack of dynamic control of transport networks by means of automated SDN-based multi-layer resource orchestration;
- Reduce energy consumption by bypassing the grooming layer;
- Foster the emergence of open industry with an open source and vendor-agnostic approach.