Application-specific policy-driven 5G Transport with Hybrid ICN

Mauro Sardara, Jacques Samain, Jordan Augé, Giovanna Carofiglio
Cisco Systems, France
Email: {msardara, jsamain, jordan.auge, gcarofig}@cisco.com

***Abstract*—The future landscape of a heterogeneous and unified access (WiFi and Cellular), where users and applications interact with heterogeneous multi-cloud networks, with many different services collaborating together, poses significant challenges to ISP, Enterprises, Cloud Providers and the applications themselves, in particular in terms of Policies, Security and QoE.**

In this demonstration we showcase the potential of Hybrid ICN (hICN) in the context of enforcing inter-domain policies: exploiting the connection-less, app-aware, multipoint transport of hICN we are able to manage policies on a per-application basis, combining multiple objectives of different players (ISP, Enterprise, Users) in a dynamic and seamless manner.

***Index Terms*—Hybrid ICN; Policy; Telemetry; 5G**

I. Introduction

The forthcoming arrival of 5G promises a paradigm shift including very high carrier frequencies with massive bandwidths, extreme base station and device densities and unprecedented numbers of antennas. Unlike the previous four generations, it will also be highly integrative, tying any new 5G air interface and spectrum together with LTE and WiFi to provide universal high-rate coverage and a seamless user experience.

Opportunities like multi-homing and seamless mobility will not be unusual anymore, becoming the norm and not the exception. Communication Service Providers will need to deal with a new dynamic and mobile set of users, each one of them exploiting different services, possibly hosted in different third party clouds. The requirements for each service differ in terms of performance, reliability, security, and policies.

The need of fully exploiting the new possibilities provided by the 5G evolution for meeting the objectives of the applications is evident. However, in this new 5G landscape other actors such as Cloud Providers, Enterprises and ISP are going to have contrasting objectives.

Applications may wish to fully exploit all the possibilities offered by the new 5G access, for instance using the WiFi and LTE channel together for maximizing the bandwidth. Concurrently, Enterprises may forbid an application from using a public WiFi access for security reasons, and at the same time Cloud Providers may force a real time application to use a specific channel because of the lower end-to-end delay. On the other hand, ISP may be interested in ensuring QoS constraints to specific services, due to commercial agreement with cloud providers (e.g. limit the end-to-end delay for a real time application).

What is required is a *policy-driven transport*, able to take into account all the requirements of each one of these actors and apply them dynamically and on a per-application basis:

- *Cloud providers* will need to manage the services end-to-end, combining different inter-domain policies and adapting them to the current condition of the users.
- *Enterprises* will need to enforce and dynamically distribute their security policies for their mobile workers.
- *ISP* may need to ensure per-application QoS constraints to Cloud Providers.
- *Applications* will need to be instrumented for both fully exploiting all the benefits provided by the 5G evolution and for taking into account policies coming from different entities, such as the Service Providers or the Enterprise, potentially in a real-time manner.

Current solutions such as MPTCP allow to exploit a multi-homed (or multi-cloud) connection in presence of a static access, but they are not able to quickly react to policy changes or mobility events, which may require applications to quickly migrate the connection, with consequent connectivity disruption and end-to-end delay variations.

In addition, the intermediate network infrastructure is not application-aware: ISP cannot provide a per-application QoS handling. This because the network/transport is application unaware: it is easy to identify a flow, but it is not easy to classify it in real time (in particular if the traffic is encrypted).

Hybrid Information Centric Networking (hICN) [1] is a relatively novel network architecture which enables a simplified and more efficient user-to-content communication. It allows to strictly tie the network and application layer, providing an application-aware network [2] which is able to discern application flows at fine granularity. In addition, it enables seamless and dynamic use of multiple networks, pull-based transport controlled at user side and seamless mobility across multiple network accesses.

In this demonstration we showcase how hICN allows to encode Application, Cloud Provider and Enterprise policies directly into the forwarding plane, allowing a more reactive and effective policy enforcement in a per-application basis. We also show how the application-aware network layer of hICN allows ISP and applications to fine customize the QoS policies depending on the service being used. Finally, we demonstrate the capabilities of hICN in quickly adapting to network changes due to dynamic policy enforcement.

II. Application-specific policy-driven transport

In hICN, application semantic is no longer decoupled by the network layer, but it is rather mapped into it. For instance, the hICN prefix `b001::aaaa::/64` may be the prefix of a real time audio stream, and `b001::bbbb::/64` could be used for the video stream.

Network policies can be enforced directly on the prefixes: an ISP could forbid the use of a particular prefix over a specific channel for QoS traffic optimization.

The available network adjacencies on a given system are abstracted as hICN faces, each one with its associated TAGS; examples of tags can be WIRED, WIRELESS, CELLULAR, REALTIME, MAX_THROUGHPUT, UNSECURE. Tags are not static, they can be updated dynamically depending on the current network conditions (a poor WiFi connection will lose its REALTIME tag if it cannot meet the end-to-end delay requirement anymore).

What a policy does is to express a preference (`require, prefer, avoid, prohibit`) on a given tag of a face. If the face does not meet the requirement, it is not selected for the forwarding.

Hybrid ICN allows to apply the policy enforced by the different actors directly on the forwarding plane: each Forwarding Information Base (FIB) entry contains information regarding the face to use for forwarding a given packet and the policy for the prefix the packet belongs to. If the network capabilities of the face do meet the requirements of the policy, the packet is forwarded. Otherwise the hICN forwarder is not allowed to use that face for sending that packet. Note that the policy enforcement directly at the forwarding plane is not possible to do within an IP network, due to the unawareness of the network layer with respect to the application.

Policies and TAGS are pushed into the FIB through an agent (fig. 1) running on the hICN nodes. The agent receives policies coming from different entities: the application itself, the ISP, the Cloud Provider and the Enterprise. Policies are delivered through the hICN control plane. Depending on the priorities, policies on one entity will be applied in place of the one specified by another. For instance, a face which usage has been forbidden by the Enterprise cannot be used, even if the Application wishes to.

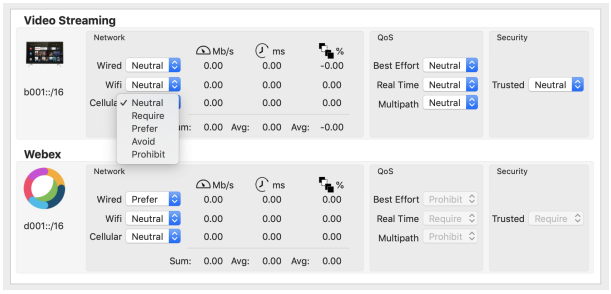

Fig. 1. Policies Agent

III. Demonstration

The demonstration consists in one remote device using two different applications with different requirements. The first is a Video On Demand application (e.g. Youtube), while the second is a Videoconferencing application.

The device is connected to an Enterprise network using three different accesses: (1) Ethernet, (2) Enterprise WiFi and (3) Company VPN over Public LTE, through which it can also access the Public Internet.

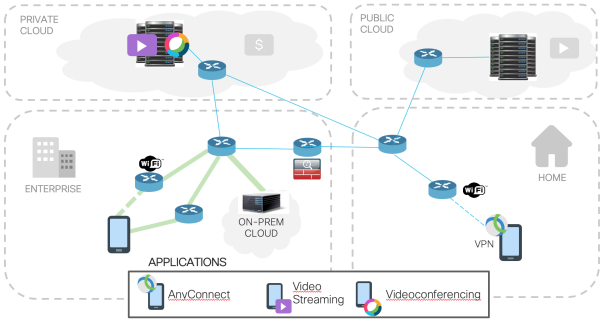

Fig. 2. Demo Topology

The services used by the applications are hosted in three different locations: (1) Company On-Prem Cloud, (2) Company private cloud and (3) Public cloud service. In the demo scenarii, the policies can be enforced by the User itself (for better showing) or by the Enterprise, which will have a highest priority with respect to the user choices.

In the first scenario we show how the user can dynamically enforce policies for maximizing the bandwidth in presence of multiple network accesses and multiple sources (the video is hosted both on-prem and on the private cloud of the company): we show how it is possible to seamlessly switch from one interface to another by changing the policy for the VoD application. We also demonstrate the hICN possibility to use the two accesses (WiFi and Ethernet) and to download from multiple sources at the same time.

The second scenario shows how the Enterprise can enforce policies through the agent, preventing the Videoconferencing application from using a public WiFi network access and forcing it to use a SECURE link, which is a VPN over LTE. We will underline how the policy choice in the agent is limited by the policies enforced by the Enterprise.

References

[1] L. Muscariello *et al.*, "Hybrid Information-Centric Networking," Internet-Draft draft-muscariello-intarea-hicn-00, Internet Engineering Task Force, June 2018. Work in Progress.

[2] M. Sardara *et al.*, "A transport layer and socket api for (h)icn: Design, implementation and performance analysis," in *Proceedings of the 5th ACM Conference on Information-Centric Networking*, '18, 2018.