COntent Mediator architecture for content-aware nETworks

European Seventh Framework Project FP7-2010-ICT-248784-STREP

Deliverable D2.1
Business Models and System Requirements for the COMET System

The COMET Consortium

Telefónica Investigación y Desarrollo, TID, Spain
University College London, UCL, United Kingdom
University of Surrey, UniS, United Kingdom
PrimeTel PLC, PRIMETEL, Cyprus
Warsaw University of Technology, WUT, Poland
Intracom SA Telecom Solutions, INTRACOM TELECOM, Greece

For more information on this document or the COMET project, please contact:

Mr. Francisco Javier Ramón Salguero
Telefónica Investigación y Desarrollo, fjrs@tid.es
Emilio Vargas 6
28043 Madrid
Spain

Phone: +34 91 337 4675
E-mail: comet@comet-project.org

Document Control

Title: Business Models and System Requirements for the COMET System

Type: Public

Editor(s): Gerardo García de Blas

E-mail: ggdb@tid.es

Author(s): Gerardo García de Blas, Francisco Javier Ramón Salguero (TID)

Wei Koong Chai, Ioannis Psaras, David Griffin, Marinos Charalambides, George Pavlou (UCL)

Ning Wang, Lei Liang, Michael Howarth (UniS)

Dimitrios Giannakopoulos, Eleftheria Hadjioannou, Chrysovalanto Kousetti (PRIMETEL)

Andrzej Beben, Jaroslaw Sliwinski, Wojciech Burakowski (WUT)

Spiros Spirou, Sergios Soursos, Vasiliki Kamariari, George Petropoulos (INTRACOM TELECOM)

Doc ID: D2.1-v1.0.docx

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
V0.1	June 2nd , 2010	García de Blas, Gerardo	Table of Contents
V0.2	July 16th, 2010	García de Blas, Gerardo	First integrated version
V0.3	July 26th, 2010	García de Blas, Gerardo	Changes in chapters 2, 3 and 4
V0.4	July 27th, 2010	Wang, Ning	General review. Changes in sections 2.3 and 2.4
V0.5	July 28th, 2010	Koong Chai, Wei	General review. Changes in chapters 2 and 3 and review of chapter 5
V0.6	July 29th, 2010	García de Blas, Gerardo	Second integrated version. All chapters reviewed. New structure of chapter 5
V0.7	August 2nd, 2010	García de Blas, Gerardo	Updates of section 4.2
V0.8	August 3rd, 2010	García de Blas, Gerardo	Updated sections 3.2, 3.4.2 and Chapter 5
V0.9	August 4th , 2010	García de Blas, Gerardo	Pre-Final version of the deliverable
V0.10	August 5th , 2010	Wang, Ning	Final review
V1.0	August 6th , 2010	García de Blas, Gerardo	Final version of the deliverable

Table of Contents

1 Executive Summary 7

2 Rationale behind content mediation 10

2.1 Introduction 10

2.2 Current roles involved in content publication and distribution 11

2.3 Problems with current Over-the-top content distribution systems in the Internet 13

2.3.1 Provider acting as distributor 13

2.3.2 Content Delivery Networks 15

2.3.3 P2P networks 19

2.4 Problem statement 22

2.4.1 Network Operators 22

2.4.2 Content Consumers 23

2.4.3 Content Creators and Content Providers 23

2.5 Benefits from content mediation and awareness 23

2.5.1 Network Operators 23

2.5.2 Content Consumers 24

2.5.3 Content Creators and Content Providers 25

3 Overview of the COMET approach 27

3.1 Objective of the COMET project 27

3.2 Basic Entities 27

3.3 Content-based Operations 28

3.3.1 Content Publication 29

3.3.2 Content Consumption 29

3.4 COMET Two-plane Approach 30

3.4.1 Content Mediation Plane (CMP) 30

3.4.2 Content Forwarding Plane (CFP) 31

4 Use cases of the COMET system 32

4.1 Use case 1: Adaptable and efficient content distribution 32

4.1.1 Rationale 32

4.1.2 Description 32

4.1.3 Benefits 34

4.2 Use case 2: Handover of content delivery path in a multi-homing scenario 35

4.2.1 Rationale 35

4.2.2 Description 35

4.2.3 Benefits 36

4.3 Use case 3: Webinar “All about CDNs” 37
4.3.1 Rationale 37
4.3.2 Description 37
4.3.3 Benefits 41
4.4 Use case 4: P2P offloading 42
4.4.1 Rationale 42
4.4.2 Description 42
4.4.3 Benefits 44
5 Business models with the COMET system 45
5.1 Introduction 45
5.2 Reference Business Models 45
5.2.1 Free content access 45
5.2.2 Charged content access 46
5.3 Business Models built over the Free Content Access Reference Model 48
5.3.1 Content Provider acting as Content Distributor 48
5.3.2 Content Delivery Networks 50
5.3.3 P2P Networks 53
5.4 Business Models built over the Charged Content Access Reference Model 55
5.4.1 Internet TV broadcasting services 56
5.4.2 IPTV managed services 56
5.5 Impact of COMET in data interconnection business models 58
5.5.1 Impact of COMET in QoS interconnection 58
5.5.2 Impact of COMET in multicast interconnection 60
5.6 Other business considerations in COMET 66
5.6.1 Impact of COMET in content publication 66
5.6.2 Interface to third-party search engines 67
6 Requirements for the COMET system 69
6.1 Introduction 69
6.2 System requirements 69
6.2.1 Global requirements 69
6.2.2 Requirements for the Content Consumers (and Content clients) 70
6.2.3 Requirements for the Content Providers (and Content servers) 70
6.2.4 Requirements for the CMP (mediation layer requirements) 71
6.2.5 Requirements for the CFP (network layer requirements) 71
6.2.6 Summary table 71
7 Conclusions 73
8 References 75

9 Abbreviations **77**

10 Acknowledgements **79**

Annex A: COMET terminology **80**

(This page is left blank intentionally.)

1 Executive Summary

The deliverable *D2.1 Business Models and System Requirements with the COMET system* is the first deliverable of the COMET project and is intended to orientate the design of the COMET architecture as well as focus the work to be carried along the project. It comprises the use cases of the COMET system, the identification of the business models that could appear on top, and the specification of the high-level requirements for the COMET System.

The document introduces the current situation in content distribution and the potential impact that the COMET system might have in that market. In particular, two main problems are identified in that ecosystem:

- The increasing number of intermediaries makes many contents accessible only for particular user communities, resulting in **global content search and direct access being fragmented**. The key issue lies in the lack of a global content naming scheme to access the content, which forces end-users to search the content through relevant intermediaries, maintaining a multiplicity of accounts, front-ends, tools and applications for content discovery and consumption.
- **Today's networks are unaware of the content they are transporting**. Due to this unawareness, networks cannot apply the most appropriate end-to-end transport strategy to provide the adequate quality of experience for the end users. Besides, flash crowds in live events which can potentially lead to traffic peaks cannot be efficiently addressed. Moreover, even when the networks are well prepared, intermediaries acting as *Content Distributors* (Internet Content Providers as YouTube, CDN providers as Akamai, or P2P platforms as Octoshape) cannot be aware of the network capabilities, traffic conditions, or the transmission requirements for the content. Therefore, the content is delivered far from the most efficient way.

These problems highlight the lack of a link between two different "worlds" (*contents* and *data transmission*), without an effective coordination between them.

The COMET system provides an appropriate linkage between both worlds, allowing *Internet Service Providers* to act as mediators for content publication and distribution. This mediation will make possible an improvement of content delivery in terms of quality and effective bandwidth utilisation, while offering unified interfaces for content access and publication, where content is treated as a *first citizen* in the Internet.

In this respect, the COMET project will follow a 2-plane approach for that mediation, including a Content Mediation Plane (CMP), in charge of offering those unified interfaces to *Content Consumers* and *Content Providers*, and a Content Forwarding Plane (CFP), in charge of the delivery of the content based on its knowledge of both the network and server status.

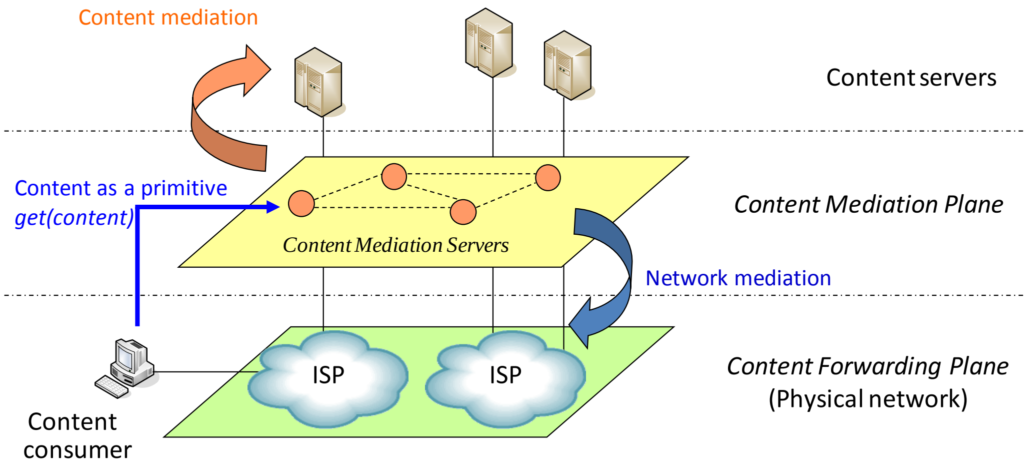

Figure 1: The COMET two-plane approach

Thanks to the mediation provided by the COMET system, key **technical advantages** can be achieved:

- Unified access to the content whatever its nature and location
- Content delivery with guaranteed QoS
- Point-to-multipoint content delivery capabilities, reducing bandwidth needs for live contents
- Graceful handover of the content delivery path, providing more resilience and flexibility for multi-homed users
- Advanced publication mechanisms, allowing *Content Providers* to update content servers on-the-fly, while switching among different ways of distribution.

These advantages are analysed in **four use cases** which will drive the design of the COMET architecture, as well as the future demonstration activities. The use cases show through storylines how the capabilities of the COMET system could be exploited in real scenarios and benefit the different actors involved in content distribution:

- ***Adaptable and efficient content distribution***, which presents the distribution through the Internet of a live event with QoS guarantees and an efficient use of network resources.
- ***Handover of content delivery path in a multi-homing scenario***, that shows the distribution without disruption of a VoD content to a multi-homed user while switching between different networks interfaces/access networks.
- ***Webinar: All about CDNs***, that covers an Internet communication service such as a Webinar, enhanced by the QoS and point-to-multipoint capabilities available thanks to the network mediation.
- ***P2P offloading***, that presents the capabilities for *Content Providers* to perform an offloading between different ways of distribution (from unicast to P2P distribution) depending on the servers and network conditions.

Through the study of these use cases and their benefits for the different actors, the business models that could emerge from the adoption of the COMET system have been identified. At a top-level, **three reference business models** are distinguished, as shown in Figure 2:

- Scenarios based on **free content access,** where content will be accessed free of charge (on the left).
- Scenarios based on **charged content access** where *Content Consumers* would pay ***ISPs*** for the contents accessed through COMET (in the middle)
- Scenarios based on **charged content access** where *Content Consumers* would pay ***Content Providers*** for the contents accessed through COMET (on the right)

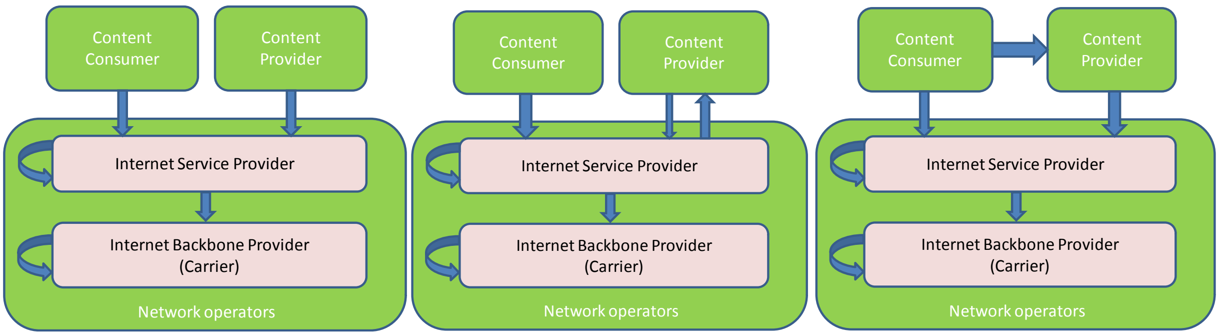

Figure 2: Reference business models

The deliverable illustrates the possibilities for the **COMET system** to encompass former and new business models related to content distribution, **becoming a flexible umbrella for all the variety of content-related business that are either present in today's Internet or those which would be potentially possible in an evolved Internet, with advanced content-**

handling features. Specifically, thanks to the content mediation provided by the COMET system, *ISPs* can be aware of the consumers' content requests. Due to this content awareness, a whole new set of business models could appear based on charging per content.

The COMET project, although does not aim at defining new business models related to Layer-3 interactions, will be a driver for *Network Operators* to settle QoS and multicast interconnection agreements and boost the development of a QoS- and multicast-aware Internet.

Finally, as a result of the detailed analysis on use cases and business models, the deliverable identifies **27 high-level requirements for the COMET system**. These high-level requirements are **currently driving the design of the COMET architecture**.

2 Rationale behind content mediation

2.1 Introduction

In recent years there has been a growing proliferation of user-generated Internet content, including blogs, photos, video, etc. The increasing trend of users generating their own content has led to an abundance of intermediaries: content is published through user websites, social networks such as Facebook or MySpace, photo sharing sites such as Picasa and Flickr, pre-recorded media aggregators such as YouTube or GoogleVideo, content delivery networks such as Akamai or Limelight, or through P2P overlay networks such as BitTorrent or eMule.

In the near future, massive content generation will not only come from end users. Companies like Blockbuster, Netflix and Apple are already providing movies through the Internet. Live content is also expected to explode once the network capabilities allow it. Moreover, journalism is focusing on small and flexible outlets with live video blogs, as an exciting business opportunity.

Finally, services like Google Street View will allow people to add live data in the future, opening the door to new business opportunities (security, entertainment, real estate or retailing are just some examples). Network and service providers struggling for extra sources of revenue may also gain from implementing presence services, for example video-conferencing and immersive tele-presence, with users tied to an identity but, because of mobility and privacy for example, not to a location.

The main problem with the current approach to content access is that there are an increasing number of intermediaries and, as a result, a lot of content tends to be accessible only by particular user communities, with **global content search and direct access being fragmented**. Specifically, the key issue lies in the lack of global content naming scheme and infrastructure to access the content, which forces end-users to search the content through the relevant intermediary, maintaining a multitude of accounts, front-ends, tools and applications in order to discover, access and consume content.

Moreover, with the current approach, the access to content needs to be machine and application-dependent. Content access requires knowing the server hosting the content and the application/session protocol used to distribute that content. This prevents content mobility, as well as the adaptation of the type of distribution in a transparent way for the end-user. A global naming scheme providing content naming persistence would bring naturally such a content mobility and adaptation, making end-users unaware of any changes in content location or application protocols.

On the other hand, today's networks are unaware of the content they are transporting. Due to this unawareness, networks cannot apply the most appropriate end-to-end transport strategy for the content in order to achieve the best quality of experience for the end users. Moreover, crowded live events that can potentially lead to network congestion peak times cannot be efficiently transported. Nevertheless, even if the networks were well prepared, intermediaries usually are not aware of the network capabilities, traffic conditions, as well as the content transfer requirements. Therefore, the content is delivered far from the most efficient way.

The same lack of infrastructure that we have introduced previously, also forces end-users to publish the content they generate through intermediaries in order to reach global audience since they do not have enough bandwidth to deliver them in an efficient and reliable way. There are multiple possibilities of achieving this but they all have some drawbacks for the publisher.

Content Delivery Networks (CDNs) implement techniques to provide server load balancing and content location independence[1], but their services are not globally available to all Content Creators, especially single users which cannot afford this kind of services.

[1] CDNs offer load balancing through dedicated servers delegating content requests to the appropriate servers. Besides, CDNs make use of dynamic request routing through the DNS, which allows changing dynamically the content servers

P2P networks are globally available to all Content Creators who want to publish their contents and distribute them through the Internet. The end-users that want to request these contents just need to have the correct application installed on their devices to download them from the source. But this type of distribution is based on best effort techniques and the protocols involved are not effective enough to support QoS (Quality of Service) for the end-user. As a result, it is difficult to plan the P2P network to provide a particular level of service to the end-user. Moreover, the perceived QoS by the end-user highly depends on network and peer conditions, which are hard to predict and control.

Providers acting as distributors (see section 2.3.1), in turn, offer the possibility of distributing end-user generated contents to the rest of the Internet but they suffer from a lack of control of the contents by the creator, that is, these creators just upload their contents to the Over-the-top Content Provider and miss the capability of having anything to do with these contents as modifying their characteristics or establishing a particular QoS for their delivery. In addition, their techniques are based on overlay distribution systems that are not aware of the network topology and cannot provide QoS to the end-users.

Given the expected exponential increase in content generation and the problems that have been previously explained, a different architecture for content location, access and distribution is necessary, providing unified content access through a new interface based on **content as a primitive**.

Such interface will make **the content both location and hosting application independent**. This new content-centric architecture would bring other advantages derived from the use of content as a primitive:

- **Ease to apply QoS or multicast techniques**. Since end-users' requests use content as a primitive, networks can become aware of the end-to-end flows related to content. That said, operator networks can transport the content end-to-end with QoS awareness. Additionally, multicast technologies and path diversity can be applied on a content-basis, typically across multiple autonomous ISP networks, which can lead to a reduction in network traffic and a more efficient use of network resources.
- **Increase of security in content delivery**. The fact that end-users' requests use content as a primitive simplifies the deployment of security policies such as content filtering and location hiding. Denying access to inappropriate content (racism, violence, etc.) becomes much simpler with this content-centric approach. Besides, content servers' location could be hidden to end users, and, in that way, servers would become inherently more protected against attacks.

The next sections provide a more detailed explanation of the problem statement with current content distribution systems through the Internet and the benefits that could be achieved with content mediation and awareness. Besides, the next sections also present a general overview of current agents involved in content publication and distribution and the systems they use.

2.2 Current roles involved in content publication and distribution

In this section, we present all the possible roles that may be involved in the content publication and distribution process. These are the Content Creator, the Content Provider, the Content Distributor, the Network Operator and the Content Consumer.

The **Content Creator** is the entity that owns the rights of the content and wants to publish it to the Internet. A Content Creator could be the creator and the owner of the content or the entity that has acquired the content and its rights. A Content Creator could be either a single end-user creating a file with their personal means (e.g. personal PC, camera, recorder, etc) or a large organization whose content is its intellectual property (e.g. music companies, movie studios). Often, the Content Creator is the same entity as the Content Provider (defined below), since it might also own the infrastructure to provide content to Content Consumers.

The **Content Provider** is the entity that specializes in storing and making content available to the Content Consumers. Content Providers (e.g. YouTube, Apple iTunes Store, etc.) are generally large companies that provide content either free or with a fee.

One of the main objectives of Content Providers is the wide availability of their content in order to reach a broad market. Additionally, they need to protect the content rights, keep the cost of content distribution low and achieve reliability and high QoS in order to provide better experience of service to their customers (consumers).

The **Content Distributor** is the entity that owns and maintains infrastructure that supports the distribution of content with adequate QoS. Content Distributors are CDNs (e.g., Akamai) that deploy servers with replicated content at various points in the network, P2P networks (e.g., BitTorrent, Octoshape) that make use of end-user resources for distributing content, or even Content Providers that maintain a distribution infrastructure. When no CDNs or P2P networks are involved, which is the commonest case, the Content Provider and Distributor are considered as a single entity.

The **Network Operator** is the entity that provides networking services, wired or wireless. The Network Operators can be classified depending on their role in the process of content distribution:

- Internet Backbone Providers (IBPs) interconnect with each other and form the backbone of the Internet. They have large capacity networks and their main responsibility is the delivery of the content to other IBPs or ISPs (wholesale services).
- ISPs have some network of their own, usually country-wide. They may serve a number of individual customers (home users or companies). Their main responsibility is the delivery of the content to the Content Consumers (retail services).

Network Operators mainly serve content related services which could become revenue sources (e.g. deployment of IPTV/VoD services). In order to increase their market share they need to provide different services or different QoS from competing operators and decrease capital and operational expenditure for efficient service deployment and maintenance.

The **Content Consumer** is the entity that consumes (e.g., watches, listens to, or reads) the content. Most of the time it is the end-user who downloads the content to their device for instant or offline consumption. In the recent Internet terminology, the term **prosumer** has also become popular as it captures the fact that an end-user can be at the same time a Content Consumer and a Content Provider.

Content Consumers need a reliable and value-for-money service with good quality in order to access the desired content depending on the characteristics of their specific device. Additionally, Content Consumers need to consume content in an efficient and transparent manner without having any concern about the location of the content, the appropriate application in order to access it or even the characteristics of the device they use. In addition, they need to be informed about alternative options of the same content.

Figure 3 represents the possible roles on the content delivery chain and their interaction in order that the content reaches the consumer.

Figure 3: Content flow from creator to consumer

2.3 Problems with current Over-the-top content distribution systems in the Internet

Over-the-top (OTT) refers to the content distribution in which providers have capabilities to go directly to consumers with their contents, bypassing traditional network gatekeepers and access providers, that is, not allowing the Network Operators they go through to take an active role or obtain any revenue from the content distribution. These contents are offered to the end-user through an unmanaged and public Internet connection.

These systems are also changing how consumers access to broadcast entertainment, and how content and communications companies provide it. They have become one of the main sources of traffic in the Internet with popular services such as YouTube, Metacafe or Google Video.

This section introduces three systems that can be currently used by Over-the-top systems to distribute their contents through the Internet. The first case is when the provider acts as a distributor, that is, the Content Provider has its own distribution system; the second case are Content Delivery Networks (CDNs), and the third case are Peer-to-peer (P2P) platforms. After some technical description of each, the different actors belonging to these systems are identified and mapped to the roles we have previously defined. Finally, the problems and drawbacks that each of these systems have are introduced.

2.3.1 Provider acting as distributor

2.3.1.1 Introduction

This section deals with the case in which Content Providers deploy their own system to distribute their contents to the Content Consumers. In fact, they play two roles in COMET terminology, as providers and as distributors.

Regarding the distribution technologies which can be used, they range from the traditional client/server architecture, in which the Content Provider has its own servers and hosts the contents inside them to be accessible by the consumers, to the more complex CDN-like approaches in which the Content Providers actually deploy a network of servers in the Internet to improve the service they offer to the consumers, as it happens with Google.

With these alternative ways of distribution, the Content Provider can get some additional advantages as cost reductions or improved QoS. One of the main functionalities they offer to the Content Creators is that they can act as mediators between them and the Content Consumers. Due to this mediation, they can offer them the possibility of publishing their contents ignoring the complexity associated with their distribution. This advantage is also present in other Over-the-top solutions as CDNs, which are explained later.

2.3.1.2 Technical description

When the provider hosting the video makes use of the traditional client/server architecture, this architecture consists of a server that hosts the video and streams it to the consumer through its Internet access connection. The consumer just needs to have a device connected to the Internet running an appropriate application to play the content (in most cases, an Internet web browser with an appropriate plug-in might be enough, as it happens with FLV videos).

Some of these services are able to deliver high-quality or enhanced contents but they need as well end-user devices capable of receiving these contents (e.g. HD video contents, interactive services, parallel Web services…)

2.3.1.3 Actors perspective and mapping to roles

The core actors, that may sometimes represent various roles, are the following:

- **Content Creator:** The author of the content. This actor can be, for instance, an amateur Content Creator (in that case it is called *prosumer*) or a large organization. When not having their own infrastructure, the Content Creators can benefit from these third party Over-the-top platforms by delegating their contents publication and distribution. These platforms provide them not only the means to distribute content to a huge base of end users but also allow Content Creators to ignore the complexity associated to the distribution of the contents. Moreover, the network traffic related to their content servers is reduced, and, consequently, fees to ISPs are reduced as well. Platforms such as YouTube, where *prosumers* can upload their videos and benefit from the services they offer, are usually free platforms.

On the other hand, when Content Creators own the appropriate infrastructure to distribute their contents, they have to publish these contents by themselves. This can also be considered as an Over-the-top solution and may be as simple as some servers hosting the videos they offer and a web portal for the users to access the contents.

- **Provider acting as distributor:** The actor that holds and makes this data accessible to others. It also operates the data centre facilities used for generic hosting and co-location. It can just consist of a server in which the Content Creator stores its contents.
- **Network Operators (IBPs and ISPs):** They provide the wide area transport for ISPs and Internet access to the end-users. They also provide Internet access to the Over-the-top platforms.

  One benefit that they can take from Over-the-top platforms that also act as distributors is that these platforms usually place their servers strategically, being aware in some way of the network topology and conditions. By doing this, the load on interconnection and the congestion in links are decreased, thus delivery and capacity costs for the Network Operator are lowered as well.
- **Content Consumers:** The people requesting the contents. They can access video contents without paying fees to the Network Operators for them, just with the Internet access these operators provide them. In some cases they must pay the providers for the contents (e.g. live sports, such as Champions League match shown on the ITV website) and in other cases they can access the contents for free, as it happens with YouTube. In this case, in which the Over-the-top platform acts as distributor to be able to offer an improved QoS in terms of delay and throughput to the Content Consumers, they can benefit from that without paying more to the ISPs.
  When acting as *prosumers*, they can benefit from an easy and free way of distributing and publishing their contents, achieving an appreciable QoS using these Over-the-top platforms.

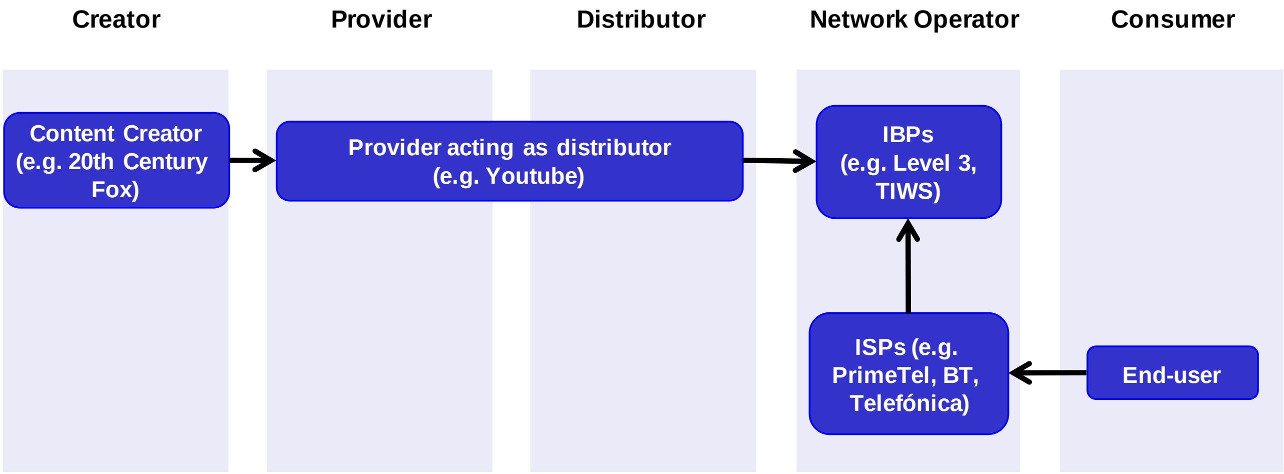

Figure 4: Actors and mapping to roles for Providers acting as Distributors

A possible mapping of the actors to roles is shown in Figure 4: . The arrows show the payment relations between the actors involved in the case where *Content Creators* delegate the storage, publication and distribution of their contents to a *Content Provider* platform.

The relation between the *Content Creator* and the *Content Provider* platform can be of different natures. Sometimes, *Content Creators* pay the *Content Provider* platform for the services it offers. The creators delegate the storage, publication and distribution to these platforms and must pay for that service. It could also appear the reverse situation, in which *Content Provider* platforms pay Content Creators for their contents in order to make them available to their clients and take benefit from that that (e.g. via advertisement revenues). In the case shown in Figure 4: , there are not payment relations between the creator and the *Content Provider* platform because YouTube is a free service and the creator just uploads the content to the platform and does not have to pay anything for the service.

ISPs charge end-users for accessing the Internet while *Content Provider* platforms would pay *ISPs* for being accessible from the Internet. In some situations, specific *Content Provider* platforms that handle a huge amount of traffic, such as Google/YouTube, could settle traffic peering (with no charge) agreements with *ISPs/IBPs* to exchange their traffic.

2.3.1.4 State of deployment

The shortcomings of IPTV have encouraged traditional broadcasting companies to go Over-the-top with video distribution today. On the one hand, IPTV services are not fully built on the potential advantages of the Internet. Instead of leveraging the power of an open environment, Network Operators with IPTV solutions keep using closed systems. The use of proprietary or “walled garden”[2] technologies and devices has slowed a widespread take-up, preventing the same open cooperation and access which people now take for granted on their PCs with the newest Over-the-top solutions.

That could be the reason why direct-to-consumer Over-the-top IPTV alternatives have emerged. Companies such as Hulu, YouTube, Apple TV, Google Video, Veoh or Metacafe have become very popular platforms to access video contents, distributing their contents with their own servers and technologies.

2.3.1.5 Problems with providers acting as distributors

Over-the-top providers distribute their contents using best-effort techniques in the Internet. These systems are not aware of the underlying network topology and traffic conditions in the end-to-end path, so they are not able to guarantee QoS to the end-user. What is more, for the same reason, they cannot provide a particular grade of service to the Content Creator itself.

Moreover, these solutions do not allow the creator to have control over the contents. The creator just uploads its contents to the Over-the-top provider and misses the capability of having anything to do with the contents. Additionally, creators have to adapt their contents to the format and characteristics that the Over-the-top provider supports.

2.3.2 Content Delivery Networks

2.3.2.1 Introduction

A CDN is a networked system of computers interconnected which cooperate —in a transparent way to the users— in the efficient and predictable distribution of contents based on geographic location, the origin content and the content delivery server. This is one of the solutions to the growing problems of distributing these contents effectively over the Internet, as the bandwidth requirements are increasing and the Quality of Experience (QoE) requirements are becoming more stringent. CDNs create copies of the content (either on demand or pre-emptively) to strategically

[2] The term “walled garden” refers here to the control and restrictions that ISPs exert over applications and content on their platforms.

placed servers hosted by CDN providers, called surrogate servers, and deliver it to end-users based on their characteristics. CDN providers are either commercial (i.e. Akamai, EdgeCast, Limelight) or academic/free (i.e. Coral, CoDeeN), and sign contracts with Content Providers. Thus, when a user requests some content from a Content Provider, the requests are automatically redirected to the closest/more adequate CDN server. The scheme of the CDN structure is shown in Figure 5.

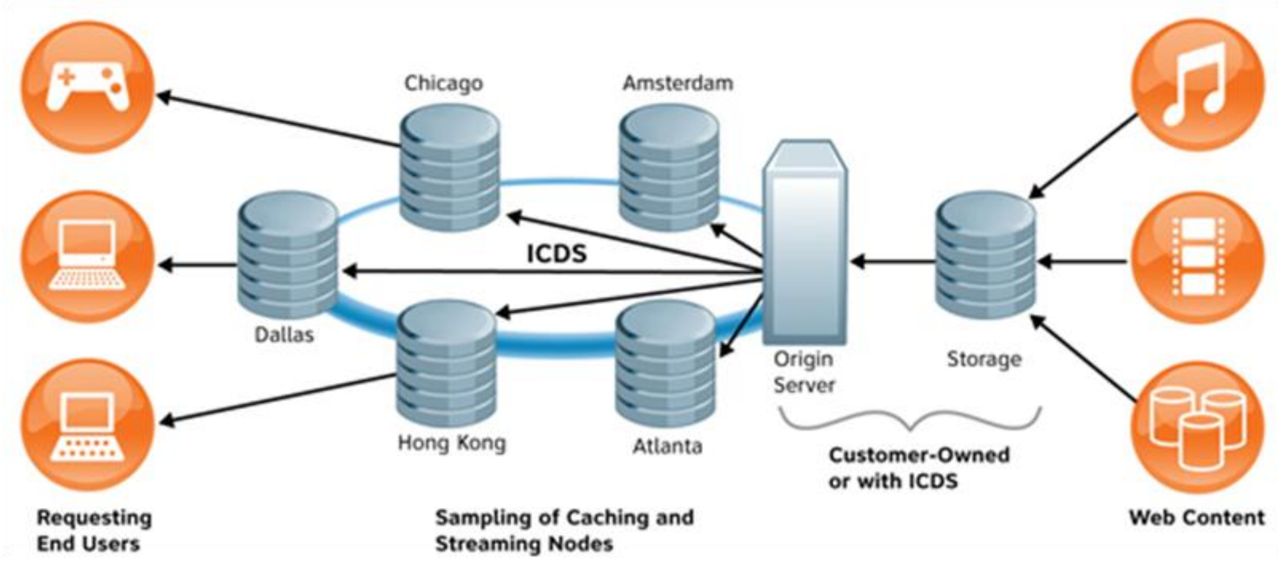

Figure 5: ATT CDN scheme (ICDS: Intelligent Content Distribution Service) [11]

2.3.2.2 Technical description

The CDN architecture consists of a **content delivery infrastructure** to host multiple copies of content, a **routing infrastructure** so that users' requests are redirected from the Web to CDN servers and copies of content are kept up-to-date, a **content distribution infrastructure** responsible for moving sources of content close to requesting users, and an **accounting infrastructure** which accounts users' accesses to CDNs and records CDN servers' traffic and usage, in order to support CDN provider's billing operations.

In general, there are two overall architectures that a CDN provider needs to choose from in order to structure their system; the *active network approach* and *the overlay approach*. The former requires the use of network components (i.e. routers and switches), which are not only used to forward packets but also, with the use of special software, to identify application types and use custom policies for each to route the content. The later uses application-specific caches and servers to distribute a set of specific content types (static web pages, live TV etc). The active network approach actively uses network components to assist the delivery of content whereas the overlay approach does not actively make use of them other than for providing basic network connectivity. Some CDN providers may use the two approaches in combination [8].

Another important aspect of the CDN infrastructure is the way that servers are placed geographically to optimize content delivery. The placement is decided using algorithms to calculate the most efficient locations [6].

Once the placement is decided, CDNs make use of content replication and caching techniques in order to distribute the content closer to the potential end users, rather than repeatedly transmitting identical versions of the content from an origin server. Since surrogate servers are not meant to serve as exact replicas of the origin server, due to their capacity limitations, only a subset of the content is usually distributed to them. Different distribution approaches are used for different circumstances and they can affect the efficiency of the end-to-end content transfer [5]. The three most popular ones are:

- *Cooperative push-based*: With this approach, the content is pre-fetched from the origin server to its surrogates while surrogate servers cooperate to reduce replication costs. The CDN holds a mapping of which content exists on which surrogate and can route requests accordingly to the geographically closest and capable surrogate.
- *Non-cooperative pull-based*: This approach is the currently most popular approach in the industry (i.e. Akamai, Mirror Image) since it does not use surrogate server cooperation which is still at its experimental stages. In this approach, requests are routed to their geographically closest surrogate server. If there is no cache entry for the requested content (cache miss) then the surrogate server retrieves and caches the content from the origin server and then serves the end user request.
- *Cooperative pull-based*: This approach combines notions from the two previous approaches. In the Cooperative pull-based approach the surrogate servers only retrieve content on a cache miss but they do so by requesting the content by other nearby surrogate servers that hold the content. They are aware of such servers by holding a copy of an index showing which surrogate servers hold the content they need [5].

Since the users require a global identifier to request each piece of content, CDNs need a mechanism to route their request to the closest and appropriate surrogate server which can more efficiently serve them. This redirection can be done either via DNS, where DNS servers redirect users' requests to optimal CDN servers, based on users' characteristics (users' location, network topology and condition, servers' health and load) or by URL rewriting, where requested URL links are rewritten and requests are redirected to optimal CDN servers. With these techniques, CDNs improve the load of the servers and the latency of the distributed content consumption.

Contents distributed by CDNs could be of three different types:

- Static, such as web pages, images, videos on demand, etc.
- Quasi-dynamic, such as continuously updated web pages, etc.
- Live content, such as streaming media, gaming, etc.

2.3.2.3 Actors perspective and mapping to roles

The CDN technology was developed to solve the problems that arise when trying to achieve efficient traffic distribution. Due to the variability of content and the different needs of each situation, different business models may apply in order to satisfy these needs. Despite the different circumstances, there are some commonalities between these models in the type of actors involved and the overall business goals [23]:

- **Scalability.** The ability to handle larger amounts of data, users and transactions. Furthermore, it needs to be scalable in a way that it can allow dynamic provisioning with high quality content delivery at low operational cost.
- **Security.** Only authorised parties must be allowed to access and modify the content. There are needs to be secure at multiple levels: physical, network, software, data, and procedural.
- **Reliability, Responsiveness and Performance.** The service needs to have high availability and able to handle outages and satisfy user experience expectations. It needs to be able to tolerate faults with appropriate load balancing.

The core actors, that may sometimes represent various roles, are the following:

- **Content Creator:** The author of the content. Sometimes they need the services of a media aggregator that provides the storage of their contents. When no media aggregator is needed, the Content Creator directly hires the services of the CDN provider to distribute their contents. In this case, the benefits explained below are also applicable to the Content Creator.

- **Media aggregator:** An entity that is in charge of storing contents, created by Content Creators. CDNs provide not only the means to distribute content to a huge base of end users maintaining a reasonable quality of experience for them, but also allow these media aggregators or Content Providers to ignore the complexity associated to the distribution of the contents. Moreover, the network traffic related to their content servers is reduced, and consequently fees to ISPs are reduced as well. Note, though, that not all the amount is actually saved, since these Media aggregators or Content Providers pay some fees back to the CDN for serving their content.

- **CDN Provider:** The actor that holds and makes this data accessible to others. It also operates the data centre facilities used for generic hosting and collocation.
- **Network Operators (IBPs and ISPs):** They provide the wide area transport for ISPs and Internet access to the end-users. They can benefit from the fact that CDN providers place their servers strategically. Usually, the CDN providers request the ISPs to locate these servers inside ISP's network infrastructure. Both cooperate to find the optimum network locations to place CDN servers, but the final decision belongs to the ISP provider. By doing this, load on interconnection and congestion in links may decrease, thus delivery and capacity costs for these Network Operators may be reduced as well.
- **Content Consumers:** The people requesting the content. With content being delivered through CDNs, the end-users have a better Quality of Experience (QoE) and QoS compared to trying to access the same content directly from the origin server, which might not be geographically aware and is more likely to have more load than a CDN surrogate server. With Content Delivery Networks, they can experience a better QoS in terms of delay and throughput —especially important in content streaming and gaming applications— than with other distribution systems. Besides, they can experiment this improved QoS without paying more to the ISPs.

A possible mapping of the actors to roles is shown in Figure 6: . The arrows show the payment relations between the actors:

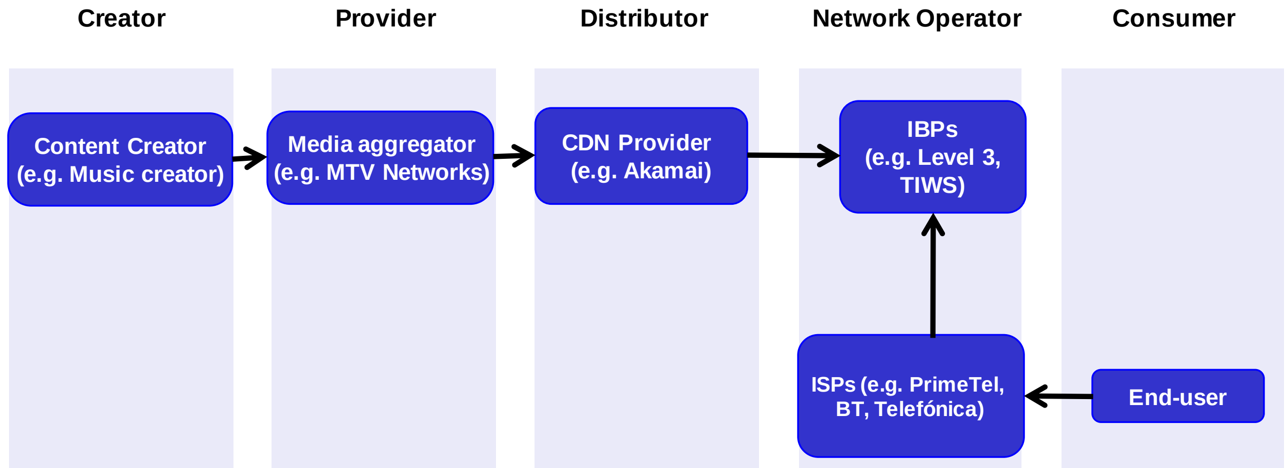

Figure 6: Actors and mapping to roles for CDNs

While in a network without CDN providers, both Content Consumers and Content Providers (in this case, this role is represented by the actor media aggregator) would pay their ISPs and Carriers for either consuming or publishing content and Carriers (IBPs) would charge ISPs for connecting them to core Internet, with the presence of CDN providers, business relationships between all actors are significantly affected.

Figure 6 illustrates the payment relations that are involved in case Content Providers delegate the distribution of their contents to a CDN provider.

ISPs charge Content Consumers for accessing the Internet, Content Providers pay CDN Providers for distributing their content, while CDN Providers would pay ISPs or IBPs, depending if CDNs are accessible by the Internet through one or the other. Carriers consider CDNs as any other network with outbound characteristics, thus an adequate traffic agreement between these two entities is arranged.

2.3.2.4 State of deployment

In the last years streaming traffic has become more and more important in detriment of P2P traffic. In fact, currently around 30% of the content that is passed through operator networks is streaming content [16]. Moreover, plenty of very popular web sites use CDNs to distribute their contents, such as CNN or RTVE.

In the past 3 years, the number of CDNs coming to the market jumped from about a dozen to more than 50 at its peak [12] and combined they have raised almost half a billion dollars while at the same time the CDN market has shown a remarkable growth [13].

An illustrative example is the CDN provider Akamai. This CDN delivers around 15-20% of the whole world web traffic through its CDN [30]. In addition, Akamai has around 48.000 servers distributed all around the world in 70 countries, and the daily web traffic it delivers is bigger than the traffic transported by many Carriers, occasionally reaching over 2 Tbps. Given such massive traffic volume, CDNs sometimes settle traffic peering (with no charge) interconnection agreements with Carriers.

Finally, some network operators, such as Tata and ATT, have entered in the Content Distribution market by building their own CDN platforms, reusing their already distributed network facilities.

2.3.2.5 Problems with CDNs content distribution

A primary disadvantage of CDNs is the increased latency in content access upon requests from Content Consumers caused by increased DNS lookup times due to the way DNS requests are routed in the DNS tree. CDN administrators are taking steps towards resolving this but it still remains a problem [7]. Another significant disadvantage of using a CDN to deliver content is the cost. It is a costly approach for multiple reasons [7]:

- Bandwidth costs
- Variation of traffic distribution
- Size of the content delivered
- Number of surrogate servers maintained
- Reliability, stability and security requirements

2.3.3 P2P networks

2.3.3.1 Introduction

The Peer-to-peer (P2P) paradigm covers a heterogeneous group of techniques, technologies and applications. This kind of networks are distributed networks consisting of participants (peers) that make their resources (such as processing power, disk storage or network bandwidth) directly available to other peers, without the need for mediation by a centralized server. This model is different from the client/server model since nodes in a P2P network work both as clients and servers, consuming and providing content to other peers on the network. Several applications can be realized over P2P systems, such as file sharing, Video on Demand (VoD) or VoIP.

2.3.3.2 Technical description

The P2P architecture allows anonymous peers to share their resources, with or without interacting with a centralized server. Opposed to the traditional client-server architecture, in P2P systems there is no concept of a content server and all participants are ‘equal’ peers. A peer gives some resources, and in return obtains other resources (including information, processing resources, request forwarding etc.) that are essential for the operation of the system and beneficial for all the peers. The peers have to deal with limited and unreliable connectivity, likely an independent addressing system and should be able to share the role of the server. Peers are autonomous and, thus, they cannot trust each other by default. In addition, P2P networks provide means to harness the power of vast amounts of computing resources, including storage, CPU cycles and connectivity from the increasing number of personal computers distributed around the network.

P2P networks can be classified as structured and unstructured. In structured P2P networks, network topology is controlled and content queries are appropriately handled by pre-assigned overlay entities with specific organization rules. In this model peers communicate with a central server (indexer) in order to publish the content they provide. Upon a request of a peer to the central server the peer receives the content from the best peer(s), taking into consideration the availability of the peer(s) offering the content and other metrics specified by the user. The most common type of structured P2P networks is the Distributed Hash Table (DHT), resulting in the efficient location of rare content with bounded searching complexity. However, structured P2P networks can become overloaded of popular data, contrary to unstructured P2P networks. The main characteristic of structured P2P networks is that peers can join or leave network in a random manner and can locate content through queries that flood the network. This model requires messaging overhead, as the request can be distributed in different levels of connected peers and the search for non popular data is likely to be unsuccessful as there is a limitation in the number of flooding steps. Lately, a hybrid form of P2P networks has appeared where servers act as mediators between peers, supporting the content delivery and collecting data about peers.

The main characteristics of P2P systems are summarized below:

- There are no distinct clients or servers, but peers assume both roles.
- Resources (e.g., content) reside on the peers.
- Peers are autonomous and independent.
- Peers and content have unique identifiers, proprietary to each P2P system.
- Infrastructure nodes (e.g., indexes) might exist to assist peers discovery.

2.3.3.3 Actors perspective and mapping to roles

The core actors that are involved in the P2P networks perspective, the payment relations and the content distribution process are presented below:

- **Content Creator:** The author of the content. In some cases it delegates the storage and publication of its contents to another entity which represents the role of Content Provider. In the example shown in Figure 7, the actor involved is a broadcasting company. When the Content Creator also acts as Content Provider, it uses directly the services of the P2P platform and the benefits and perspectives explained below for the case of the Broadcasting Company are also applicable.

  As *prosumers*, P2P network users can benefit from the low barriers to entry when not aiming at large audiences with high availability or grade of service.

- **Broadcasting Company:** This entity, which is not a mandatory actor in P2P networks, can help Content Creators with the storage and publication of their contents.

  As concurrent media distribution to large audiences might be relatively expensive with traditional ways of distribution (e.g., CDNs), Content Providers sometimes use P2P platforms. Theoretically, P2P distribution could reduce this cost, so providers could benefit from delegating content distribution to P2P platforms instead of maintaining their own hosting services or using traditional CDN services to reach wider audiences with cost efficiency.

- **P2P platform:** The actor that provides the P2P infrastructure. The evolution of these P2P platforms has shown a trend to add additional servers to contribute to the distribution, migrating from pure P2P infrastructure to a less decentralized way of distributing their contents, closer to a traditional CDN (e.g., Joost).

- **Network Operators (IBPs and ISPs):** They provide the wide area transport for ISPs and Internet access to the end-user. P2P traffic percentage is still large today (about 50%) and varies significantly across different ISPs, but it is generally declining in percentage [16]. The distribution cost of this traffic is shifted to the ISP, as peers become servers and upload data instead of downloading. In addition, the egress traffic of the ISPs is increased and no revenue is added from serving the content.

What is more, P2P overlays are created over ISP networks without any awareness of ISP topologies, boundaries and policies. This usually results in higher interconnection costs and inefficient use of network resources for the ISPs. From an ISP point of view, a P2P network distribution does not add value and probably increases the cost, because of increased bandwidth usage and high interconnection traffic. For this reason, sometimes it is interesting to have sources of P2P traffic inside the ISP network to save interconnection costs and partially recover the costs induced for the internal traffic distribution. Following this interest, ISPs are studying network-aware P2P solutions to reduce costs (inter- and intra-domain) by exploiting locality by protocol modifications (IETF ALTO WG [17], EU FP7 SmoothIT [18]) or by local caching.

- **Content Consumers:** The people requesting the contents. As consumers, P2P end-users can get, in case of congested servers, better QoS from their Content Providers if content is distributed by P2P. Besides, the low barriers to entry of this way of distribution for *prosumers* help in terms of content availability and, accordingly, end-users can have access to a vast amount of niche content, rarely available in other kinds of distribution infrastructures.

A possible mapping of the players to roles is shown in Figure 7. The arrows show the payment relations between the actors:

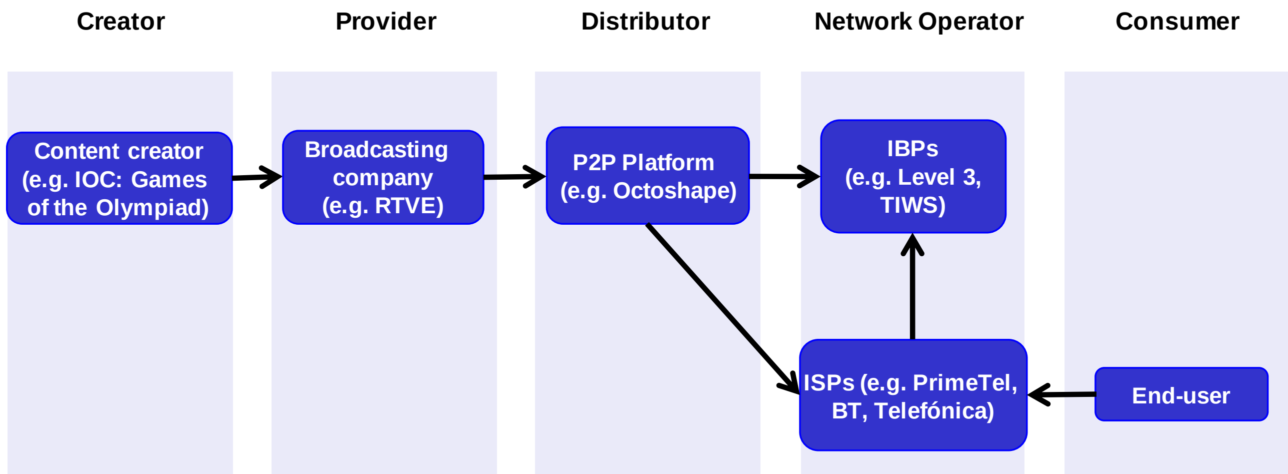

Figure 7: Actors and mapping to roles for P2P Networks

Content Distributors using P2P networks connect Content Providers with Content Consumers, through Network Operators in the Internet and potentially over many ISPs. The associated business relationships that are shown in Figure 7 are explained below.

Content Providers usually pay P2P platforms for the delegation of their content distribution. Sometimes, despite P2P usually claims low cost and somewhat efficient distribution, these platforms can find the way of charging end-users for access to their exclusive content, turning them into premium users (e.g., SkypeOut). There are also P2P platforms accessible for prosumers that can delegate the distribution of their contents without paying for this service. This is the case of BitTorrent or eMule P2P networks.

ISPs could explore ways of sharing costs with P2P Content Distributors. They could also create new business models based on providing better QoS guarantees to providers/consumers, particularly in the access network.

2.3.3.4 State of deployment

P2P networks and CDNs act in a rather competitive way, as P2P Content Distributors typically ask Content Providers to pay less than what CDNs charge. P2P systems have been used for applications such as distributed computing, collaborative and content sharing software, file sharing, messaging software etc., since the underlying technology is more suitable for such applications. However, because P2P reliability is low, his success has been limited to services not aiming a guaranteed Grade of Service to the end-user.

Content Providers have tried to use P2P platforms to distribute their content in order to save costs migrating from traditional solutions (CDNs, hosting…). But, as far as they have been unable to provide enhanced grades of service to consumers, more controlled and easier to manage ways of content distribution have become a preferred solution for most providers.

As mentioned before, Joost have migrated their systems from pure P2P to a less decentralized way of distributing their contents. They originally used only P2P streaming technology to distribute content. Later on, infrastructure nodes (stable servers) were introduced as the original seeders of all content and to improve the availability of rare content. Eventually, in late 2008, Joost switched to the traditional client/server model, using a Flash-based Web player.

Currently, content distribution systems based on P2P networks are generating almost half of the global Internet traffic [16].

2.3.3.5 Problems with P2P content distribution

One of the reasons for centralising content distribution is the difficulty in planning the P2P network to provide a particular level of service to the end user. Moreover, the QoS perceived by the end user highly depends on network and peer conditions (e.g. peer churns), which are extremely hard to predict and control. QoS further degrades with the lack of control at the access network. Moreover, specific client software is required to consume contents using these P2P platforms, while usually just a web browser is required with traditional distribution methods. Finally, a key reason for the migration from P2P solutions to more centralized ones is the fact that nowadays P2P traffic is not handled well by the ISPs, due to the negative effects in planning that it usually introduces.

2.4 Problem statement

Content-based networks have been proposed lately to address several problems that the current Internet is facing due to its host-centric nature. In particular, it has become apparent that while users are searching for content or services, the Internet, as a host-centric entity, points to the host or machine where the content is located rather than to the content itself. In addition, the increased complexity of today's networks and the unprecedented growth of user-generated content pose significant challenges to the operational efficiency of today's content resolution algorithms.

Furthermore, members of specific user-communities only can reach some content, making it difficult and sometimes impossible for other interested users to reach the desired content. There is common consent by now that there is no unified global content access and resolution architecture. Below, we list the problems and requirements that have been identified for different roles of the COMET system, namely, Network Operators, Content Consumers, Content Creators and Content Providers.

2.4.1 Network Operators

In recent years there has been a growing proliferation of content in the Internet, which is either generated by single users, big corporations or media sites. The delivery of pre-recorded or streamed content has led to an increase in consumed bandwidth, which is expected to grow in the coming years. The current host-centric approach to networking makes it impossible for routers and network entities to identify content by name. As a result, (web-) caching and replication techniques are difficult to implement, since predefined Points of Presence (PoP) have to be decided and setup ahead of time.

Moreover, content delivery in the current Internet is offered without QoS guarantees and resilience. Over-provisioning is the current common practice among operators to guarantee enough bandwidth for the end-user. In content-based networks, however, the network will be able to inherently provide QoS and guarantee that content is available wherever and whenever it is needed or is becoming popular.

Furthermore, it has been noticed in the recent years that some content are popular among people around the same geographical area (e.g., sports, national elections etc.). IP-multicast has not been widely deployed in the past, due to both weak business models and technical limitations. However, the gain of IP-multicast and similar content-dissemination techniques, where content is replicated and distributed to many simultaneously interested users, is expected to increase resource-management efficiency and enhance traffic-engineering optimization for Network Operators.

2.4.2 Content Consumers

As mentioned earlier, the host-centric Internet paradigm restricts users from accessing all available content in the Internet. The naming architecture of the current Internet, which is based on IP, points to a single machine that hosts the desired content. This, however, allows Content Providers to make their content available to groups of users only and with no option of providing the appropriate QoS. This is the case, for instance, with P2P user communities, newsgroups etc.

Furthermore, best effort users in the current Internet suffer from bad QoS with the experience becoming worse in case of multimedia applications (e.g., online TV and video) and even more so for interactive applications (e.g., voice and video over IP). Although operators try to overprovision their networks, limited revenue from such types of content distribution reduce their ability to overprovision to the extent that best-effort users can receive acceptable service quality. Efficient content naming and mediation architectures will enable operators to manage resources more efficiently and therefore, provide better than best effort services.

2.4.3 Content Creators and Content Providers

Currently, content publication and delivery poses significant barriers for smaller scale Content Creators such as single users or small corporations. A large number of single users that create their own contents, for example, have to publish their content through intermediaries (e.g., YouTube) in order to distribute their content, mainly due to the lack of the necessary network resources (e.g., bandwidth).

On the other hand, Content Providers are not currently able to guarantee QoS for content delivery in the Internet. As a result, they degrade the quality of content delivery, by reducing the resolution of video for instance, as a compromise against low network resource availability (e.g., bandwidth).

The COMET system will address these issues by designing an efficient naming and addressing architecture that is going to handle user requests by pointing to the content itself rather than to the actual machine that hosts the content. Furthermore, network aware routing and resource allocation protocols will guarantee, inherently, QoS to end-users, better business options and opportunities for operators and a unified publication procedure for Content Providers.

2.5 Benefits from content mediation and awareness

In this section, we list the benefits that the COMET system is going to provide. As in the previous section, we differentiate between the main roles: i) Network Operators, ii) Content Consumers, and iii) Content Creators/Content Providers.

2.5.1 Network Operators

Operators will be able to reduce their bandwidth demand by efficient resource allocation, especially at the domain edges and within their own backbone. In that sense, operators will reduce costs from infrastructure development, due to the capability of content server selection, e.g. an ISP could select a server located inside its network instead of a server in other ISP domains, and by using different distribution schemes such as multicast or peer-to-peer.

It is estimated that around 30% of the content that passes through operator networks is streaming content [16]. Since the servers of those Content Providers are distributed all over the world, international transit is needed in order to retrieve this content, making operator networks less cost-effective.

One need that has arisen for the Network Operators is the location of the surrogate servers. The optimal solution would be to have these servers within the operator's national geographical boundaries in order to save international bandwidth. It is widely known that location of surrogate servers is closely related to the content delivery process. As such, greater emphasis needs to be placed on the issue of choosing the best location for each surrogate.

Network-awareness is going to be incorporated in the COMET system design enabling operators to efficiently manage traffic within their networks and offload busy links and servers wherever and whenever this is needed. This feature will reduce the operators' cost regarding additional infrastructure development and will improve the end-users' quality of experience. The extent to which the above will happen is going to be evaluated during the course of the project. The monitoring modules that will be designed are expected to supply network providers with the required information about the conditions of their respective domains. In addition, caching techniques are going to provide sophisticated offloading of network traffic from within the operator's own domains. In particular, caching techniques for popular contents at the edges of an operator's network are expected to reduce traffic from the core of the domain.

Network Operators will benefit from the above techniques since their deployment will result in minimizing the overall network bandwidth consumption for transferring replicated content from the servers to the Content Consumers. Content Consumers will also benefit since optimal surrogate servers placement will reduce user perceived latency for accessing real-time contents (e.g. gaming applications).

New business models can be developed, where QoS and resilience in content distribution can be offered to both Content Providers and users.

In recent years, operators are focusing on providing new services, besides connectivity and bit-shifting, that could eventually become revenue sources. It is increasingly evident that those new services involve content (e.g., operators deploying IPTV/VoD "walled gardens" or carriers offering CDN-like services available in the public Internet). The possibility to offer new online services will increase the operator's revenue.

Being able to inherently provide increased QoS to all customers, the operators can develop and offer new services to users, or extend and develop further already deployed services, such as IPTV, online shopping, etc.

2.5.2 Content Consumers

Users will access multimedia content in the Internet with the same or higher quality as the traditional media.

One of the main concerns of Content Consumers is the Quality of Experience (QoE) they receive when consuming content. Therefore, QoE is the key factor for the success of the COMET system. Although QoE is a subjective measure, it is actually reflected from the combination of the Quality of Content (QoC) and the QoS. As such, improved QoC and QoS translate to higher Content Consumer satisfaction. Network QoS is itself influenced by several factors such as end to end delay, packet loss, etc., but also the time spent to serve the content (i.e. the time between a consumer request and the content being delivered).

Also, other factors playing a role in user satisfaction are, of course, security and reliability. Security needs for users are, among others, privacy and charging trustiness, while reliability needs to deal with the anytime availability of the content —in the COMET system, even if a copy of a content is moved or if a server is down, the same content could be retrieved from a different source.

Simplification of access to content

The growing proliferation of user-generated content (e.g. blogs, photos, video, etc.) through various intermediaries (e.g. different social networks, photo sharing sites, pre-recorded media aggregators, peer-to-peer networks, etc.) have created various fragments in the Internet where specific contents may only be available to subscribers of certain communities or services. Thus, Content Consumers may still have to go through multiple search operations (including possibly installation of different software or clients) in various avenues to locate a piece of content. Such complicated (especially to those who are not Internet savvy), tedious and often time-consuming content access process is a problem in today's Internet.

This highlights the need for a simplified and unified content access process that defragments the Internet and offers consumers access to all kind of multimedia contents through the Internet (one channel to reach all kind of contents). If a piece of content exists within the Internet, it should be easily located and accessed. However, it should be noted that, before accessing any content, consumers are required to have the necessary rights or permissions for that content (especially copyrighted content).

Content searching can be more powerful and comprehensive, integrating in one result list all related contents no matter if they were stored in a P2P overlay network or in the web.

Consolidation of the different communications channels – such as satellite broadcasted TV – on a single IP based channel

Users will get access potentially to all kind of multimedia contents through the Internet (one channel to reach all kind of contents), which will simplify the access to content.

Users would be able to use the applications they want or already know to access any kind of content from different locations and different connectivity channels such as mobile networks, DSL, or others, as long as they support the IP protocols stack.

2.5.3 Content Creators and Content Providers

The architecture proposed by COMET will bring the following benefits to Content Providers and Content Creators:

Content Providers would like to make content widely available, so as to reach a broad (global) market and maximize their revenues.

Content Providers, nowadays, deliver text, pictures, audio, video, applications and services to end-users. Some limitations exist where streaming applications, either live or video on demand, are concerned. Content Providers need to take advantage of the COMET architecture for bandwidth-intensive, streaming media applications since traffic from these applications will constitute the majority of Internet traffic.

Also, Content Providers would like to have control over who can access the content - by user, by country, by region- to protect the content rights.

All Content Providers will obtain bandwidth savings, and therefore cost savings, by exploiting different distribution schemes such as multicast or peer-to-peer.

Small Content Creators/prosumers will benefit from lower entry barriers in terms of cost. They will be able to reach hundreds or thousands of users from their houses to broadcast live events (radio, TV shows, etc.) without too much bandwidth costs thanks to the capability to exploit network multicast, peer to peer, server distribution, etc.

Content Providers will benefit from the quality of service provided to them by the operators in terms of reliability, reachability of users, users QoE, etc.

Operators will provide Content Providers the means to achieve the quality of experience, reliability and reachability requested by them. The COMET system will support various different technologies such as different distribution schemes, server load sharing, network bandwidth QoS enforcement adjusted to the content properties, dynamic and transparent content distribution method updates, a simple interface for the Content Provider to inform about content changes, etc. that will guarantee the Content Providers satisfaction.

3 Overview of the COMET approach

This chapter presents a high-level view of the COMET approach. Throughout the text, specific COMET terminology will be used and explained at the same time (e.g. Content Name, Content ID, Content resolution, etc.).

3.1 Objective of the COMET project

The COMET project sets out to define a content-oriented Internet that provides fast and simplistic content access with network-awareness and user-unawareness that is gradually deployable over the current and the Future Internet via the provision of a unified COMET interface. In the core, there are two major areas that require re-thinking and design:

- A unified, secure and scalable Internet-wide content organization structure (from global content naming and addressing to publication of the content to the discovery and finally the delivery of the content to Content Consumer) that defragments the various content access and distribution avenues regardless of the nature and requirements of the content (e.g. real-time vs. elastic, pre-recorded vs. live etc.).
- A content transport system supporting different types of distribution modes (unicast, inter- and intra-domain multicast, anycast, peer-to-peer) that is adaptable to the network and content server conditions.

3.2 Basic Entities

The COMET system consists of two major conceptual entities: Content Mediation Server (CMS) and Content Aware Forwarder (CAF).

- **Content Mediation Server (CMS)** – This is a key conceptual entity we introduce in this project that is responsible for content manipulating, for instance content publication, resolution and delivery operations[3]. All these processes need collaborations between CMS entities, in which case communication protocols are necessary between CMSs for fulfilling specific tasks. Although conceptually the CMS is one entity, in real implementation, it can be a combination of several physical machines with each of them providing separate service but interconnected to form a coherent mediation plane. Similar to the current DNS [14][15], we assume each participating domain or ISP will have at least one CMS entity. The number of CMS per domain will depend on performance and resilience considerations; in case of a multiple CMSs implementation within a domain, the load will be shared between them via anycast techniques, primary-secondary approaches, or others. In order to ease an incremental deployment of COMET over the Internet, it will not be mandatory to deploy CMS in all of the parties involved in the end-to-end transmission over the Internet (ISPs, transit carriers, etc.) Hence the lack of CMS in any of these segments, although it might affect the overall performance in some circumstances, it would still be valid to provide some of the most relevant COMET features, such as delivery with QoS, uniqueness of content identifier, or server selection.
- **Content aware Forwarder** - We also envision another key entity for enabling content distribution across the global Internet - content-aware forwarders. It should be noted that CMSs are not necessarily responsible for actually carrying the content from the content source to individual Content Consumers. Such a role is fundamentally fulfilled by content-aware forwarders. Compared to legacy IP routers, content aware forwarders[4] are capable of

[3] The role of CMS in content delivery is not to actually carry the content traffic, but to manipulate content delivery paths according to network conditions.

[4] Here, the use of the word forwarder is preferred over the word router to highlight that the main functionality is the forwarding of content packets.

processing content packets based on their identifiers in a native way, while it is not necessary to force all routers within an ISP's network to be content aware. In such a scenario, content aware forwarders form a virtual network infrastructure for actually delivering the content. To enable the most suitable content delivery paths in dynamic network environments, the CMS entity is responsible for providing guidelines and instructions to its local content aware forwarders in treating content flows within the network.

Besides the two major entities defined above, in order to understand, two other entities identify the physical end points involved in the end-to-end content distribution. These are the following:

- **Content Server** – The server that actually hosts content is referred to as Content Server. It should be noted that Content Servers are typically maintained by Content Providers, and hence it is not actually part of the COMET system.
- **Content Client** – The actual end host machine which is the destination of a content flow. In practice, Content Clients are owned by Content Consumers. Like the entity of content server, Content Client is not part of the COMET system.

Based on the description above, we can see Content Providers and Content Consumers are defined as roles (see section 2.2), while Content Servers and Content Clients are the physical machines owned by Content Providers and Content Consumers respectively.

3.3 Content-based Operations

Before we introduce specific content manipulation operations in the COMET system, we first define the following terms related to content:

- **Content**: the actual piece of data that the Content Consumer requests (web-page, file, video stream etc.).
- **Content ID**: an opaque (i.e. non human-readable) string that is used by COMET system in order to point to the specific content requested by the Content Consumer.
- **Content Name**: a human-readable string that explicitly and uniquely identifies the requested Content.
- **Metadata**: a set of parameters related to a specific content which are given to COMET system and identify the content properties required for accessing a piece of content (e.g. content servers, ways of distribution, QoS requirements). Besides, these metadata can include a set of keywords that could be used in a COMET-aware search engine.

Now we describe specific content-based operations within the COMET system. Fundamentally, we can envision two types of operations: *content publication*, which is Content Provider oriented, and *content consumption*, which is Content Consumer oriented. Figure 8 schematizes these processes and represents the involved entities (Content Servers, Content Providers, Content Clients and Content Consumers).

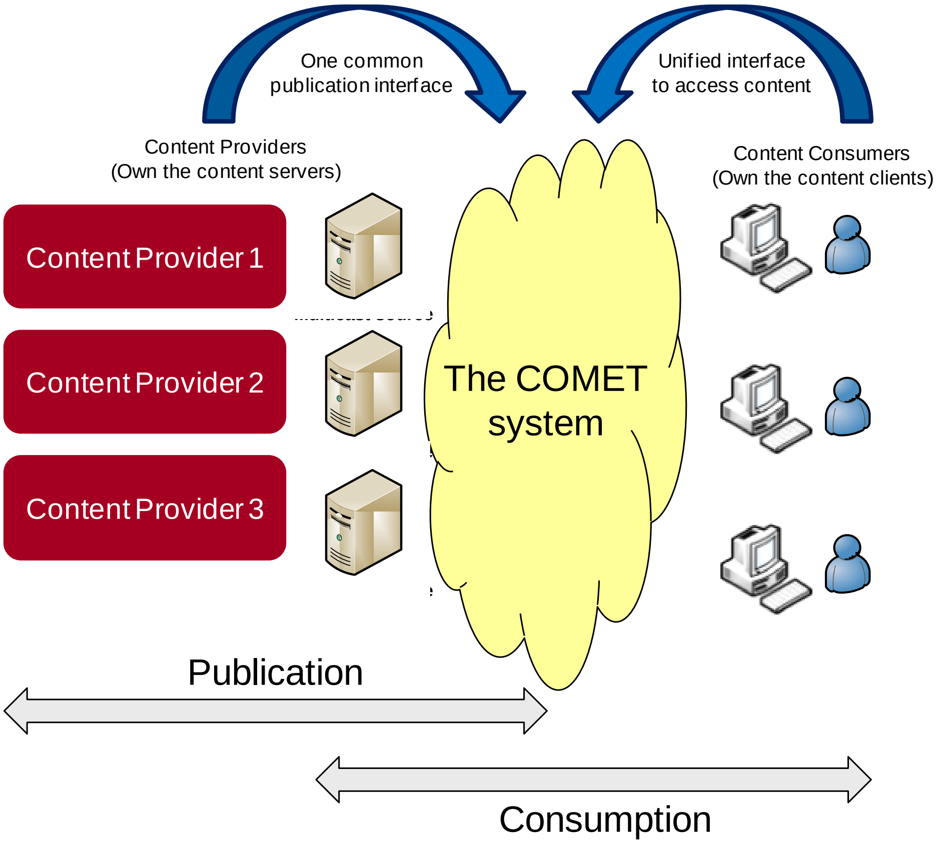

Figure 8: Schema of Content Publication and Content Consumption

3.3.1 Content Publication

Content publication is the process of making content available to Content Consumers.

Replicas of a specific content object are stored in content servers owned by Content Providers. Once a Content Provider has new content, it needs to notify the COMET system, specifically the CMS entity, for making it available for access across the Internet. This type of notification to the COMET system initiated by the content provider is called *content registration*. This registration will be made through a global and unique COMET publication interface accessible for all Content Providers. As a result of the publication process, a unique content identifier will be returned to these Providers. As mentioned before, content is physically hosted on content servers that are owned by the Content Provider.

Once a content has been registered in the COMET system, individual CMS entities are responsible for publishing it across the global Internet so that the information about the content is available throughout the Internet. This operation is called *content register dissemination* and is an optional process of the COMET system.

3.3.2 Content Consumption

Content consumption is the process initiated by a Content Consumer to receive the requested content. This process is today mostly transparent to the Content Consumer who just clicks on a hyperlink or introduces a URL in a web browser on his own machine (content client). With COMET, the consumption process will be triggered from the Content Consumers' request through a unified interface offered by the COMET system. Upon the consumer's request, (1) the COMET system first tries to locate the actual content server that hosts the requested content, and after that, (2) the content server delivers the content to the Content Consumer's device, making use of standard session and transport protocols.

The process of locating the content can be further divided into the following two stages:

- *Name resolution* – This process is responsible for translating the human-readable content name to a machine-oriented content identifier (ID). The content ID will then be used for the content search operation between CMS entities.
- *Content resolution* – The second resolution process is called content resolution. This process is responsible for discovery of the requested content based on the given content ID. In COMET, this process will be able to locate all the copies of the same content if the content has been replicated and hosted at various content servers. This is crucial for optimization / enhancement of the content delivery and also, enabling capabilities such as anycast. It must satisfy the requirement of location independence in content-oriented network (i.e. the final location of the host should not be revealed). There are different ways to decide the "best" copy. For instance, a possible option would be to use the number of hops as the metric in deciding the nearer copy. Alternative options will be considered within the context of the project, considering also QoS requirements and metrics. A more sophisticated scheme may involve the computation of the network resources in the domains involved in the transportation of the content and the current or even projected future load at the content server(s).

3.4 COMET Two-plane Approach

The COMET project will follow a 2-plane approach, including the Content Mediation Plane (CMP) and the Content Forwarding Plane (CFP). In this section we describe how the aforementioned content-based operations are actually mapped onto these two planes.

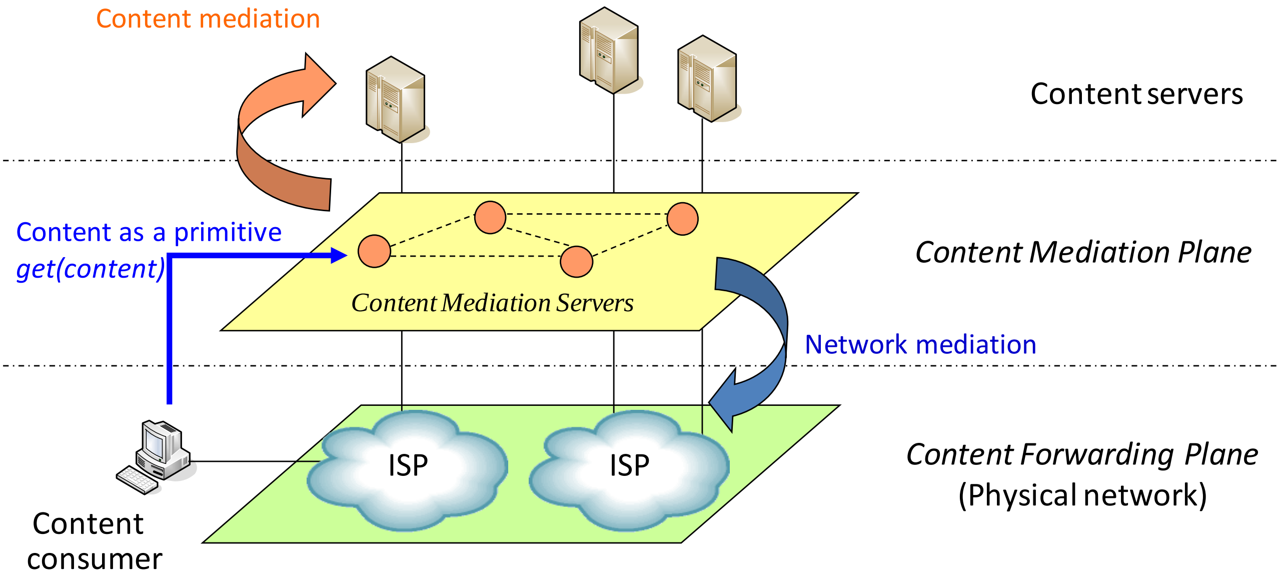

Figure 9: The COMET two-plane approach

3.4.1 Content Mediation Plane (CMP)

Fundamentally, the basic task of the CMP is to locate the actual content server that hosts the requested content. In case multiple content servers have the requested content, it is the job of the CMP to find the best candidate according to specific criteria, such as server load conditions and the associated delivery path quality. From this point of view, both the aforementioned name resolution and content resolution functions are typically enforced in the CMP.

From the description above, it is not difficult to infer that the name/ID resolution operations in the CMP are actually fulfilled by individual CMS entities. As shown in Figure 9, individual CMS elements constitute the horizontal CMP for handling content requests from Content Consumers. The CMS elements themselves may form a virtual infrastructure for locating the content server that hosts the requested content. The meaning of "mediation" is twofold as far as the CMP is concerned:

on one hand, the CMP interacts with individual content servers, e.g. probing their loading conditions such that Content Consumer requests can be intelligently resolved to the most appropriate content server in order to achieve optimized server load performances. Such a function is called "content mediation". In addition the CMP also interacts with the underlying network (represented as Content Forwarding Plane) in order to provision/select optimized content delivery paths from the selected content server back to the requesting Content Consumer. Such a function is called "network mediation".

3.4.2 Content Forwarding Plane (CFP)

As shown in Figure 9, the Content Forwarding Plane (CFP) is in charge of the *delivery* of the content once requests are resolved based on its current knowledge of both the network and server status. Thus, once the requested content has been found (or the best copy is being identified for the case where multiple copies of the same content is available), the CFP will be responsible for enforcing the actual path back to the Content Consumer from the identified content source. Although not shown in Figure 9, the role of content aware forwarders effectively fulfils the functionality of the CFP in actually delivering the content traffic in an optimised manner.

It is obvious then that the two planes (CMP and CFP) are not independent of each other and are required to communicate to achieve smooth and complete operation of content access (i.e. from request from Content Consumer to the final delivery of the content back to the requester). As an example for such a vertical interaction between the CMP and the CFP, the CMP will require information from the CFP to achieve network-awareness, while, on the reverse direction, the CFP will require information from the CMP in establishing the most suitable paths for the transportation of content with specific QoS requirements. This bi-directional interaction of both planes is elementary in achieving a seamless content delivery.

For detail technical description of the COMET architecture, please refer to *D2.2: High-level architecture of the COMET system*. *D3.1: Interim Specification of Mechanisms, Protocols and Algorithms for the Content Mediation System* will provide the in-depth description of the content naming architecture and the content and name resolution operations. Finally *D4.1: Interim Specification of Mechanisms, Protocols and Algorithms for the Content Mediation System* documents the delivery process under the COMET architecture.

4 Use cases of the COMET system

The use cases that are introduced in this chapter refer to the distribution of contents and services in which there exists mediation in the network. Given that Network Operators implementing the COMET system will be content-aware, new and different scenarios are possible regarding content distribution through the Internet.

Next sections 4.1 to 4.4 describe different use cases that have been considered as part of the COMET project and will drive the design of the COMET architecture, as well as the future demonstration activities. It must be noticed that these use cases do not describe business models, which are covered in chapter 5, although they provide the business rationale and the expected benefits behind them. They also explain through a storyline the role of the COMET system as content and network mediator. In this respect, it must be highlighted that the COMET architecture is still under design and, therefore, some details could change in the course of the project, although the main concepts of the use cases will remain valid.

4.1 Use case 1: Adaptable and efficient content distribution

4.1.1 Rationale

Currently, the distribution of live events through the Internet demands high capacity links in the Content Provider and several servers to cope with the required processing capacity. CDNs offer alternatives to provide server load balancing and content location independence, as well as their processing capacity and content replication techniques, but their services are not globally available to all Content Creators, especially single users which cannot afford this kind of services.

Even when a Content Provider is able to hire the services from a CDN, the amount of traffic in the ISP network is not reduced, since the CDN cannot help to distribute the traffic inside the ISP network. If multicast techniques were globally available in the ISP network, in the same way that happens today with other operator services (e.g. IPTV), and they were offered to third parties — such as CDN providers or directly to the Content Providers themselves— the traffic could be distributed in a more efficient fashion, reducing the server and link capacity requirements for these providers.

In addition, Content Providers cannot offer QoS guarantees for the content delivery to the end user, even when hiring CDN services[5]. Content Providers and CDN providers could benefit from QoS capabilities offered by ISPs in order to deliver their contents in a guaranteed way.

Furthermore, in current content distribution, the way to access the content depends on the content nature. For instance, once a live event is finished and is made available as VoD content, the way to access to that particular content usually changes because of the different way of distribution. If the way to access the content were independent from the content nature, for instance, with a content identifier independent from the type of distribution, the consumption process from the end user's viewpoint would be the same, thus simplifying his access to the content.

4.1.2 Description

In this use case, let us suppose that a TV corporation wants to distribute a live football match through the Internet, using the most convenient transmission modes available. During the match, the TV channel identifies that more distributed capacity is needed for streaming and it is dynamically added without disrupting with the current connections. After the match finishes, the channel continues offering it with the same identifier, but as VoD content.

[5] This fact is particularly relevant in wireless accesses, where QoS can be severely degraded unless premium traffics are specially handled.

Figure 10 depicts a schema on how the COMET system could handle this use case. It shows the different interactions between end users, Content Providers and the COMET planes (Content Mediation Plane and Content Forwarding Plane). The numbers correspond to the sequence in the storyline below.

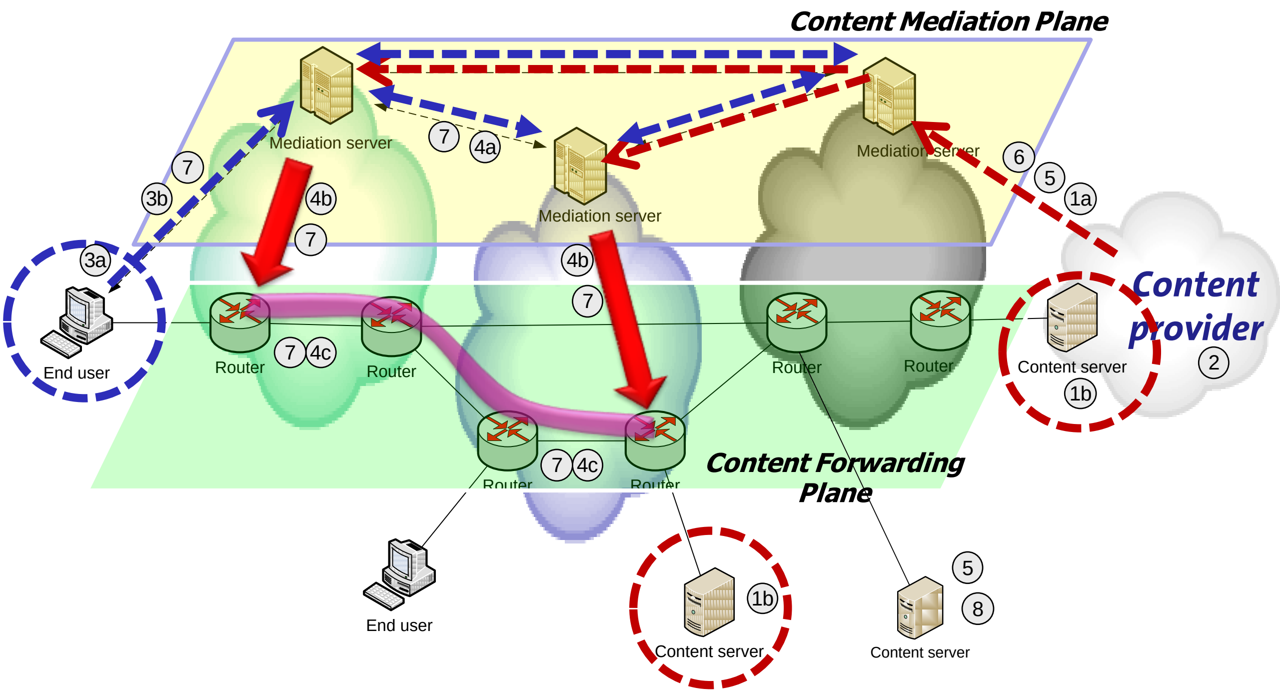

Figure 10: Schema of the use case

The storyline of this use case could be as follows:

1. The TV channel prepares for the live broadcasting of the football match. For that purpose, the channel:

   a) registers in the COMET system the details of the content stream and associated metadata (number and type of substreams, QoS requirements, etc.) and the set of live streaming servers that will be supplying the content together with the associated distribution means (anycast as there are several geographically-distributed sources, and the fact that multicast should be used where possible for live content).

   b) obtains a unique identifier (Content ID) and, optionally, a human-readable identifier for the match that could have been specified by the provider itself (Content Name).

2. The TV channel starts to broadcast the live content.

3. End users launch applications in their end devices to request the live broadcast by:

   a) obtaining the unique content identifier via either a search engine, electronic programme guide or another out-of-band mechanism (e.g. link in a blog, word of mouth, newspaper article, etc.)

   b) requesting the video stream from the *Content Mediation Plane* (which is formed by the distributed network of Content Mediation Servers) using the unique content identifier.

4. The *Content Mediation Plane* resolves the request, identifies the server to be used and triggers the establishment of the communication path according to the requirements of the content (QoS, resilience, cost, etc.)

a) This includes mapping the content identifier (and possible sub-flow IDs) to the optimal streaming server (anycast mode), based on load/performance/distance metrics (for instance, it could bind it to the closest streaming server).

b) The selection of the path for the content delivery from the server to the originating user needs to take into account the QoS, cost, resilience, etc. requirements for the particular content in question; in this example it is live content and low latency is therefore important.

c) If available in the local ISP, a multicast group can be used in order to optimize the use of the network.

5. During the live broadcast, in case more capacity is needed, new live streaming servers can be transparently added by the TV channel just by updating the registry (associated to the unique identifier) in the COMET system.
6. Once the live broadcast has finished, the TV channel updates the content registry in the COMET system to associate the identifier to the recorded version of the same event, with the details of the new VoD servers and associated meta data to show this is non-live content so multicast is not allowed, QoS requirements may be relaxed, etc.
7. End users can now connect to the VoD version of the content making a request to the *Content Mediation Plane* using the same unique identifier. This step is analogous to the case of live content with the difference that the number of sub-flows might be different. The content-aware forwarding plane is now instructed to associate the flow to a less restrictive class of service, and multicast capabilities of the local ISP may no longer be used.
8. In case of a reduction in the number of content requests, the TV channel might reduce the number of VoD servers. The content identifier will remain the same, however, so any changes in the location or number of servers will be handled by the content mediation plane and will be transparent to the users.

4.1.3 Benefits

This scenario presents the following benefits:

- The same content can always be retrieved by using the same identifier regardless of whether the event is live or recorded.
- Worldwide server capacity can be transparently adapted on-demand while the content is being served. This might be particularly useful in live events with an unexpected or difficult to predict audience.
- Content delivery is transparently associated to the most appropriate class of service/network plane, which ensures the users receive the QoE they expect and that the network is used efficiently as possible.
- Multicast capabilities can also be used to optimize the use of resources, minimizing transmission costs in the local ISP.
- The same advanced publishing mechanisms would be also available (and feasible) for small content producers.

4.2 Use case 2: Handover of content delivery path in a multi-homing scenario

4.2.1 Rationale

The explosion of multimedia services in the Internet calls for sophisticated mechanisms that are able to guarantee acceptable service quality even to best-effort users. One of the many approaches to deal with this issue is path selection according to the characteristics of the corresponding route.

With increasing popularity of multi-homed end hosts (both content servers and consumers) connected to the Internet, typically through both wired and heterogeneous wireless accesses, it is possible to exchange traffic between two hosts following multiple distinct paths, but without necessarily being supported by any multi-path routing protocol within the network, as long as at least one of them is multi-homed. In this use-case, we are going to explore issues of multi-homing, without necessarily relying on any multi-path routing protocols for adaptively performing path switching against dynamic network conditions. More specifically, multi-homed content consumers may initiate the request to perform local handover from one interface to another during a content consuming session. Typical reasons for such an interface handover may include: quality maintenance or even enhancement, preference in using different access networks (typically for mobile content consumers), etc. Such an operation might be regarded as an example of the Edge Controlled Route (ECR) paradigm to be investigated in the COMET Project, which allows higher flexibility in consuming the content at the user side.

4.2.2 Description

Multi-homing (MH) at the host level provides inherent resilience and native “multi-path selection” given that a multi-homed host has more than one network interfaces/addresses and is physically connected to more than one ISP networks (or two access points provided by one single ISP). As mentioned above, such a feature does not need to be supported by any dedicated multi-path routing protocol. Now we describe how a multi-homed content consumer is able to request a handover of the content delivery path so that the content reaches the content consumer through a different local interface. Such intelligence for requesting handovers is effectively embedded in the COMET client component at the content consumer side, which is responsible for interacting/coordinating with the core COMET functions at the Content Mediation Server (CMS) side. The storyline, shown in Figure 11, could be as follows:

1. Initially the content consumer is receiving a VoD content from one if its interfaces with a dedicated IP address.
2. During the session it may make a handover request to the CMS entity at the ISP side through the alternative interface it would like to activate.
3. Upon receiving such a request, the CMS entity decides to perform a handover of the current content delivery path to an alternative content delivery path connecting to the new interface from the content server. This decision could be based on:
    a) The preferences from the content consumer, expressed in the handover request
    b) The overall network conditions. To achieve this, individual CMS entities at the ISP side may keep up-to-date information of end-to-end path quality by active measurements and may change the content delivery path according to traffic dynamics. In fact, this approach has been proposed before in the literature as an overlay measurement technique [32].
4. As a result of the handover, the VoD content will be watched without disruption by the end-user, who will not notice the handover at all. The COMET system will follow the “make-before-break” strategy during the process. More specifically, the new content request will be first issued from the alternative interface and only after the content is delivered through the new path, the original connection will be dropped. Since such an operation is enabled at the Content Mediation Plane (CMP) below the application layer, the actual content consuming application will not even notice such changes. As such, the seamless handover operation done by the COMET system enables disruption-free path switching to the application Over the top.

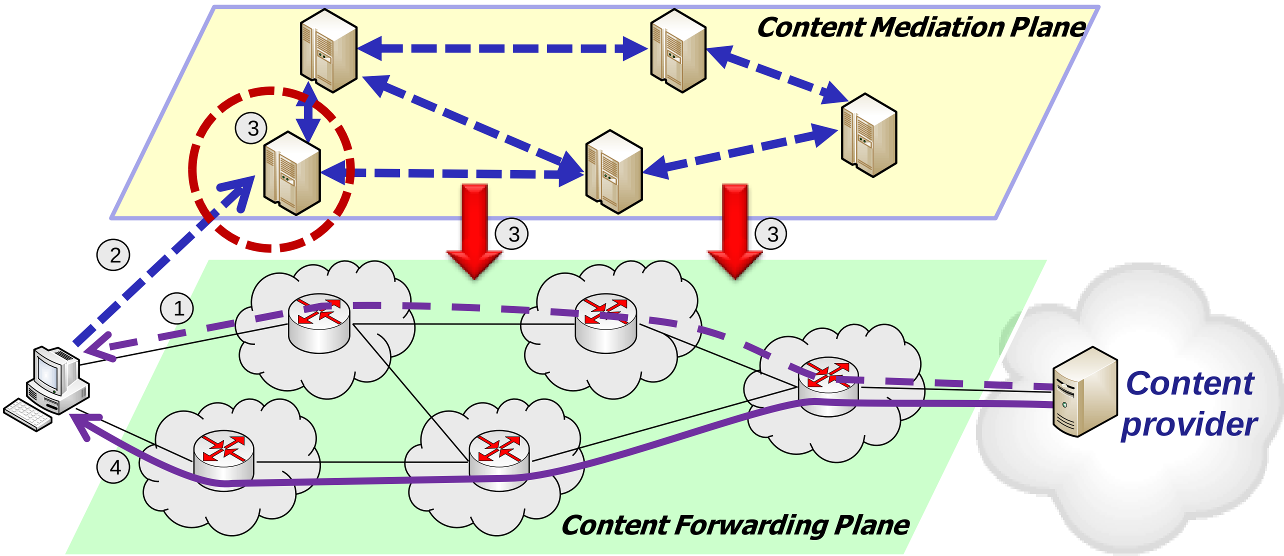

Figure 11: Handover of content delivery path in a multi-homing scenario

It has to be noted that the COMET system, as a mediator between the content handling and their actual physical delivery, is very well suited to perform this kind of transparent path switching without causing any disruption in the application layer. In fact, the COMET system is assuming that content consumers are always unaware of the physical location of the content servers, and vice-versa, thanks to the content mediation concept. This property is highly exploited in this use case by demanding to the COMET system the capability to perform the handover just changing the content delivery path while preserving the same appearance for the applications in both endpoints.

Regarding the strategies to perform the handover of the content delivery path, different alternatives can be considered. For instance, multicast strategies could be applied based on request forwarding (similar to "join request" messages in PIM-SM multicast protocol [34]) and further pruning of former delivery branches. In any case, the main idea is that a new branch of path is created and followed by the content flow based on the corresponding state maintenance. Meanwhile the original path leading towards the IP address associated with the previous interface will be pruned.

4.2.3 Benefits

This multi-homing approach is expected to have the following advantages:

- Provide enhanced network resilience for multi-homed end users. In case of route failure which may take up minutes to re-converge, the affected content consumer may immediately switch to alternative path by changing the destination address associated with local interfaces.
- Provide transparent interface switching for multi-homed end users, without any disruption on the application layer and the sessions previously established.
- Provide enhanced QoS to best-effort users in dynamic network environments, which is not the case in today's Internet.
- Provide enhanced network resource utilisation performance across domains, achieving win-win situations at both the network side and the customer side.

4.3 Use case 3: Webinar "All about CDNs"

4.3.1 Rationale

The appearance of the Internet has opened multiple new communication alternatives for people to exchange knowledge and ideas. One of these communication channels is the webinar. "Webinar can be a presentation, lecture, workshop or seminar that is transmitted over the Web" [21]. The key feature of a webinar that differentiates it from other Internet applications is that it allows interactive meetings to happen online, making online meetings better resemble physical meetings. What makes it attractive over the physical meetings is the reduced cost, the location-independence for participants, and the absence of the need of very specific equipment. The features may vary between the different implementations of a webinar but typically they involve the transmission of voice, video and data. In webinar terms, there are usually two main roles in a webinar scenario; the *presenter* and the *attendee*. The presenter is the person initiating the webinar and typically the one driving the discussion. The attendees often play the role of listeners but they can also interact, ask questions or even take the role of the presenter during the meeting. An example of a webinar implementation is the OpenMeetings software [33].

As mentioned above, the use of webinars has increased in the recent years, especially in the field of education. Many of these platforms use several web conferencing technologies and the communication between the end-user and the webinar platform is done via the Internet. This is the reason why the webinar is so popular and accessible by a large number of users.

However, the current structure of the Internet does not provide guaranteed QoS, which could create problems in the distribution of contents from the webinar server to the end-user and vice versa. Furthermore, the sending and receiving of multi-content by such a system make more relevant the need for differentiation of multiple streams. Another important problem that is inherited from current naming architectures is the inability of dynamic publishing and updating. This is a clear disadvantage for the webinar because it makes the access and churn of participants inflexible.

In this use case we will try to explore and give solutions to the problems discussed above by providing guaranteed QoS in different types of streams (Video, VoIP, etc.) and explore the possibility of dynamic publishing. In the first case, COMET will help to provide end-to-end QoS, possibly creating different paths for each type of content inside a Web-Inar (audio, video, desktop) depending on the content requirements. The flexibility of the COMET naming architecture will allow that all the contents that are part of the Web-Inar have a unique Content ID. What is more, given the features of the COMET publication process, it will be possible to dynamically update the parameters of the contents that are inside the Web-Inar.

4.3.2 Description

A webinar platform allows offering live webinars and also the possibility of recording the webinar so that it can be saved and cached on edge servers.

In the webinar architecture, the central component is the origin server to which both attendee and presenter eventually have to be connected. The webinar content flows through the origin server and is transmitted accordingly to the registered parties. There are edge servers that help with the caching, acting as proxies for the webinar content, and are necessary to provide a good and efficient service, although they are not mandatory. If a webinar setup has to work without edge servers in its infrastructure, then the origin server needs to be a machine that can handle the additional load which would have been reduced with the use of caching techniques. The audiovisual stream of the webinar will be transmitted and distributed via a webinar platform to a large number of users. During the webinar video/audio/desktop will be broadcasted among the users taking part in the webinar, who will have the opportunity to interrupt the presentation for comments or queries.

Using these capabilities of the webinar platform architecture in combination with COMET we can offer fast download speed and high content availability to the end user. The collaboration between the COMET system and the webinar platform architecture will offer real-time communication for video/audio/file transfer and desktop sharing. Moreover, knowing some of the characteristics of the network (end to end bandwidth, jitter, etc.) and also stream parameters (e.g. bit rate, video codec, etc.), the appropriate paths can be selected to enforce those parameters onto the network to transmit the stream through one of the servers of the webinar platform with the required QoS.

With respect to the webinar use case, live streaming of a football game becomes substantially different despite the fact that both the webinar and live football streaming are real time events. Thus they differ in the way they are streamed and the way the service is provided. In the webinar case there is a unicast connection between the client and the edge/origin server allowing the interaction between them e.g. to submit a question (audio/live chat) to the presenter, open to all users. Additionally, a client could make a presentation once such a request was approved by the presenter. Thus, in this case, a particular client would be changing its role, becoming the receiver stream broadcaster. This switching of roles will result in affecting the network since, in this case, the attendee of the webinar will become the presenter, temporarily changing the direction of the stream, so that the attendee becomes the broadcast sender and the presenter becomes the broadcast receiver. All of these changes are activated from the CMS in cooperation with the webinar. The Content Mediation Plane finds the new paths for each role and uses the CFP to enforce the changes in the network.

To ease the discovery as well as the transmission of the content through a webinar, a common functional identifier, named Content ID, will be used.

During the creation of the webinar, the CMP generates multiple Content IDs, one for each type of content. Each of these identifiers corresponds to a different set of metadata for each type of traffic separately (VoIP traffic, video receiver/sender, chat etc.). This set of metadata includes information for the content flow (source/destination), servers' location and also information about user requirements on Quality of Service. All this information is encapsulated in the corresponding Content ID and spread throughout the Content Mediation Servers.

The webinar scenario is a set of multiple contents for each role (presenter or attendee) meaning that a single content ID, as it is specified before, cannot describe all the set of contents. **For this case, COMET introduces the new concept of *Meta-Content ID***. The idea is to use a single Meta-Content ID to bind multiple inter-related conventional Content IDs together. For the case of the webinar, two different Meta-Content Ids, one for the presenter and one for the attendee, must be created.

Figure 12 shows how the presenter and the attendee can send or receive multi-content using the Meta-Content IDs. Thus the Meta-Content IDs include a set of Content IDs that have been specified from the webinar in cooperation with the CMP. Each of these Meta-Comet IDs is assigned a specific role. For example, MCID1 is for the role of Content Creator (presenter) and MCID2 is for the role of end-user (attendee). During the webinar, either the presenter or the attendee can change roles, which means that each of them can be assigned to a different Meta-Content ID.

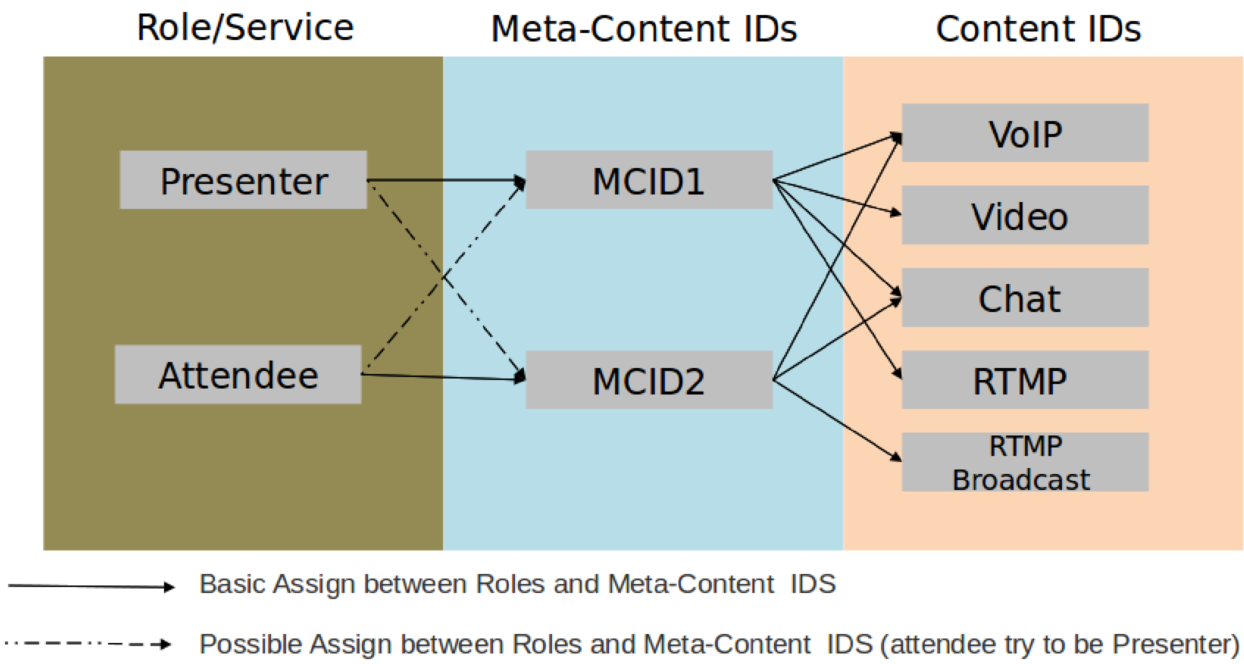

Figure 12: Meta-Content ID and webinar use case

The way that the COMET system will undertake the webinar use case is shown in Figure 13, where the different interactions between client/presenter, webinar servers and the COMET planes (Content Mediation Plane and Content Forwarding Plane) are shown.

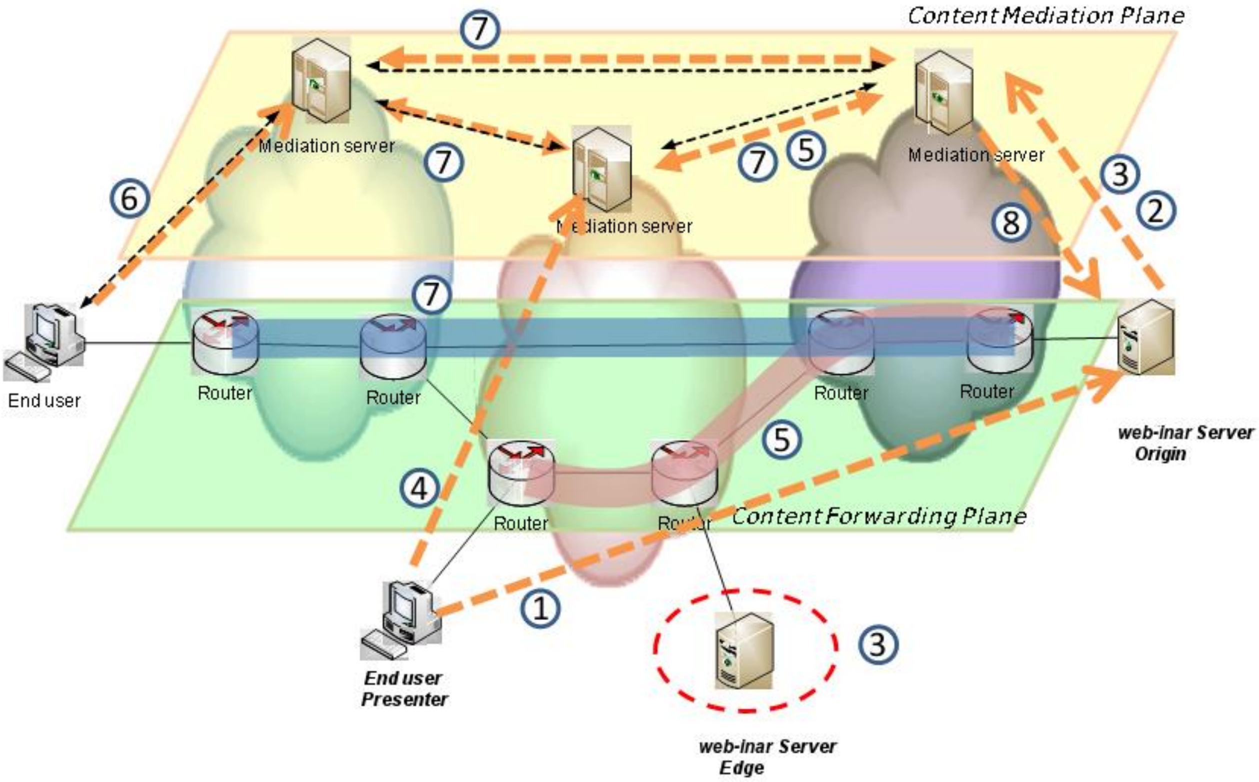

Figure 13: Simple Use Case Scenario

The storyline of this use case could be as follows (where the numbers match the ones in Figure 13):

1. The user-presenter (who acts as the Content Creator) creates a webinar, prior to the meeting, specifying the time and the duration, as well as the subject of the webinar. At the same time, on the server of the webinar platform, a Broadcast ID is created automatically.
2. The webinar obtains two unique identifiers and two human-readable identifiers, one of each for the presenter and another one for the attendees.
3. The webinar platform enters into the COMET servers the information about the different components of the Meta-Content required for the operation of the system and assigns the QoS requirements for each of them. In addition, information about the location of the edge servers and some technical characteristics about them are also entered into the COMET servers.
4. The presenter initiates the webinar procedure and finds the unique Meta-Content ID (using a search engine) according to his role. This request is received by the Content Mediation Server located in the domain of the presenter.
5. The CMP receives the presenter's request, finds the location of the origin server and tries to establish a communication channel between the presenter and the origin server according to content's demands (i.e. stream demands). When the communication channel has been established, the webinar starts.
6. One of the clients finds the Meta-Content ID for the webinar using some search engine or through published invitations. This request is received by the Content Mediation Server located in the client's domain.
7. The CMP receives the request, identifies the server of the webinar platform to be used and triggers the establishment of the communication path according to the stream requirements. The information the COMET system processes is the load of the server, the available bandwidth in the network and the types of QoS requirements for this request.
8. The CMP eventually forwards the request to the server of the webinar platform and the client is subscribed to the broadcast stream. So the user will start receiving the multi-content stream.

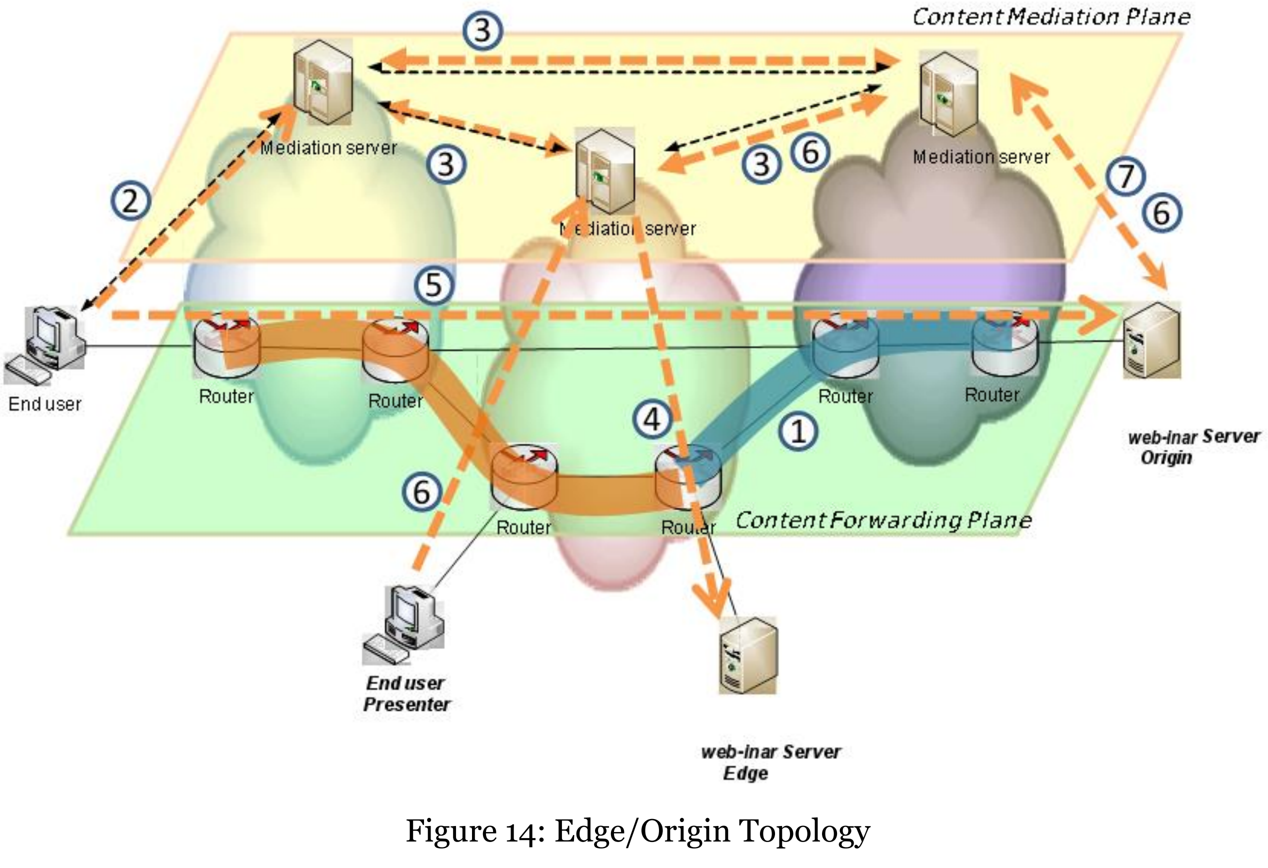

Figure 14: Edge/Origin Topology

To conclude with the description of the use case, a more complex scenario can be analysed, involving the use of origin/edge servers as well as the clients' interaction with the presenter. The steps of the storyline of this advanced scenario are the following:

1. A scenario/topology that includes an origin server and an edge server is created. The edge server is subscribed to a Broadcast ID and the origin server, via a transmission protocol, sends the webinar stream in real-time.
2. The client finds the webinar's unique identifier and requests to subscribe.
3. CMP receives the request, and checks the QoS requirements for the specific content as well as all the available servers (edge/origin). The COMET mediation plane decides that for many reasons the origin server cannot serve the client, so it finds the nearest edge server and delivers the client's request to it.
4. An edge server of the webinar platform receives the request and the client becomes subscribed to the broadcast stream.
5. During the webinar, a participant client request to pose a question about the current presentation (the interaction can be made either through VoIP or some other instant messaging but on both cases is realized through a transmission protocol and the web-browser).
6. The presenter, who has control of the webinar, based on Remote Procedure Calls (RPC) protocols, can change roles with another client, making the client temporarily the presenter and vice versa. Thus, in this case, the flow of information has changed for these two users. The client now starts sending the transmission stream from the device to the origin server and the presenter is the one subscripted to the Broadcast ID. Even though the roles change and the stream changes its source, the session remains the same, making the change transparent to all the clients connected to the webinar.
7. The origin server, after receiving the RPC requests, notifies CMP and changes the stream direction for these two users. That will last until the presenter cancels the client's rights and takes back his original role.

Finally, the webinar can be recorded and accessed by any user as a pre-recorded content in a web server.

4.3.3 Benefits

This scenario presents the following benefits:

- The same content can always be retrieved by using the same identifier regardless the event is live or recorded.
- The efficiency and QoS of the online meetings are improved. COMET will help to provide end-to-end QoS, possibly creating different paths for each type of content inside a Web-Inar (audio, video, desktop) depending on the content requirements. Since webinars are acting more and more as replacements of face-to-face meetings, this functionality is becoming more essential for the participants.
- The publication and reachability of webinars are improved. The COMET naming architecture which will uniquely identify the set of different parts of the webinar will allow these contents to be more flexibly handled. Also, the publication and updating of the different webinar parameters will be easier through the COMET architecture.

4.4 Use case 4: P2P offloading

4.4.1 Rationale

Currently, the distribution of live events through the Internet demands high capacity links in the Content Provider and several servers to cope with the needed processing capacity. But it could be the case that at peak periods or when unpredictable events occur, the provider serving capacity was not able to cope with this demand and serve all the audience. Under these conditions, the provider is bound to reject connections in order to keep the grade of service that the end users are willing to receive (even paying for it).

This situation is not desirable for the provider, which might lose some potential clients willing to receive the content even in a degraded modality. In order to avoid this situation, the Content Provider could make use of alternate ways of distribution such as P2P streaming, acceptable to provide a best-effort distribution at peak periods and more cost effective for the Content Provider. However, this situation cannot be achieved today since there is not a common interface to access the content independently from the way of distribution.

With a common interface independent of the way of distribution, a provider will not have to reject potential clients and will offer a best-effort service even in bad conditions.

4.4.2 Description

In this use case, a Content Provider wants to distribute a live event through the Internet, using a traditional client-server model. In addition, when the servers are fully busy, being unable to serve the streaming to additional viewers, the provider would want to dynamically offer the possibility of serving the content via P2P streaming, so that the content remained accessible to all potential viewers, even in a degraded fashion. In such a situation, new clients will be requested to use their P2P streaming applications to access to the content streaming. This change in the way of distribution would be made without disruption to the previously connected users and transparently for the end user thanks to the use of the same COMET identifier for all modes of distribution.

For the new clients to be able to access to the same content independently of the distribution system (either traditional server or P2P networks), it is necessary to integrate the content publication in P2P networks with the COMET system, so that content identifiers in COMET have some kind of mapping to the content identifiers used in the content platform.

Figure 15 depicts a schema on how the COMET system could undertake the use case. It shows the different interactions between end users, Content Providers and the COMET planes (Content Mediation Plane and Content Forwarding Plane). The numbers correspond to a sequence in the storyline below.

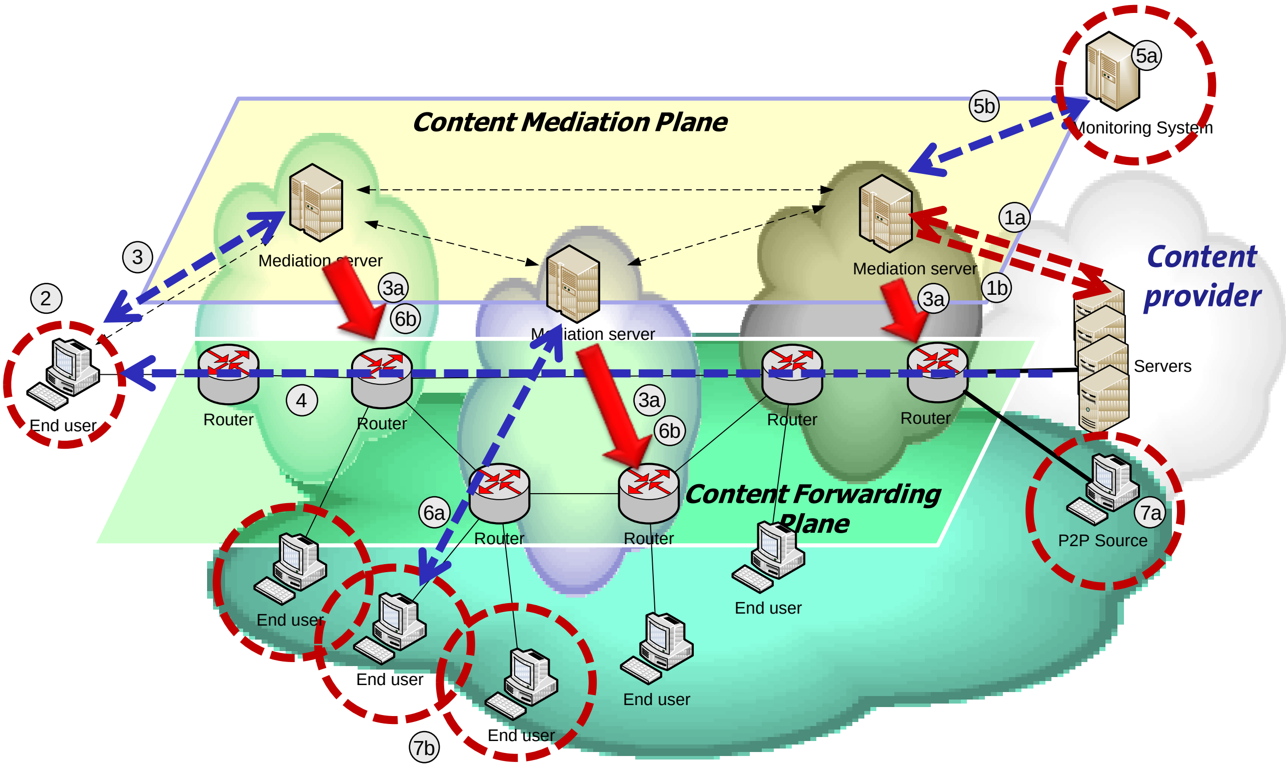

Figure 15: Schema of the use case

The storyline of this use case could be as follows:

1. The Content Provider prepares for the live broadcasting of the event. It will prepare two kinds of servers (normal streaming servers and P2P streaming sources) and will register them in the COMET system. For that purpose, the Content Provider:
    a. Registers the content in the COMET system, its two ways of distribution, the details for each way of distribution (set of sources, protocols, MIME types, etc.) and the conditions under which to switch between them, pointing out that P2P streaming should be used only in case of server or path congestion. Optionally, the TV channel specifies a Content Name for the content.
    b. Obtains one unique global identifier for the content (Content ID), valid for both ways of distribution. Additionally, it will be necessary a mapping between the COMET content identifier and a content identifier in the particular P2P platform. Either the Content Mediation Systems or some servers in the P2P platform must be aware of such mapping, although the specific details are still to be defined.
2. End users use their end devices to request the live broadcast. The end user obtains the unique Content ID or, optionally, the Content Name via search engine, electronic program guide, or out-of-band mechanisms (e.g. word of mouth).
3. The video stream is requested to the Content Mediation Plane (CMP) using the Content ID. The CMP returns the parameters that the application needs to retrieve the content from the provider.
    a. If needed, the Content Mediation Plane prepares the Content Forwarding Plane accordingly.
4. First users will retrieve the content stream directly from the servers.
5. In case the uplink capacity of the servers is full, the COMET system dynamically switches the way of distribution to P2P streaming. Here is assumed that:
    a. In the CMP, there are systems with knowledge about servers' load/BW.

   b. In case of server congestion in terms of load or bandwidth, the CMP will become aware and will change the way of distribution to P2P streaming.

6. End users trying now to consume the content with the same Content ID will be guided to use their P2P streaming applications to retrieve the content.
   a. The video stream is requested to the Content Mediation Plane (CMP) using the same Content ID.
   b. If needed, the Content Mediation Plane prepares the Content Forwarding Plane accordingly.
7. The content is distributed now via the overlay network built by the P2P streaming application:
   a. The source node provided by the Content Provider is connected to the P2P overlay network as one of the seeds of the content
   b. All new users are connected to that P2P streaming overlay network and can act as distributors of the content (as in any other P2P distribution).

4.4.3 Benefits

This scenario presents the following benefits:

- The content can always be retrieved by using the same identifier, which is agnostic of whether the type of distribution is based on client-server model or on P2P streaming.
- The content can be delivered to all the audience even when the provider does not have enough bandwidth or server capacity to serve them through a client-server model.
- The provider's method of distribution could be transparently adapted on-demand while the content is being served. This might be particularly useful in live events with an unexpected or difficult to predict audience.
- The same advanced publishing mechanisms would be also available (and feasible) for small content producers.

5 Business models with the COMET system

5.1 Introduction

This chapter presents the business models that could emerge on top of the COMET system. It is worth mentioning that the COMET framework is not strictly confined to one single standard business model but aims at becoming a flexible umbrella for all the variety of content-related business that are either present in today's Internet or those which would be potentially possible in an evolved Internet, with advanced content-handling features.

At a top-level, two kinds of scenarios can be distinguished:

- Scenarios based on **free content access** and where ISPs would shape a content-oriented Internet. Content Consumers would pay ISPs for the additional COMET capabilities for content resolution and delivery, while Content Providers would pay ISPs for the COMET capabilities for content publication and delivery.
- Scenarios based on **charged content access** where Content Consumers would pay for the content to be accessed through the COMET system. These scenarios are further classified depending on the business entity charging the content consumer, having *ISP-based content access charge* and *Content provider–based content access charge*.

This chapter is structured as follows. Section 5.2 describes the reference business models of COMET, describing the entities and their payment relations, according to both kinds of scenarios (free, non-free). Then, sections 5.3 and 5.3.3 present different business scenarios taking the reference models as basis, but incorporating some particularities. Next, Section 5.5 covers the impact of COMET in the current data interconnection models, dealing specifically with the QoS and multicast data interconnection. Finally, section 5.6 includes other important business considerations in COMET, such as the content publication through COMET and the interfaces towards third-party search engines.

5.2 Reference Business Models

5.2.1 Free content access

This reference business model covers the scenarios of free access to content in the Internet through an enhanced connectivity provided by COMET-capable ISPs and carriers. Content is published by Content Providers through the publication interface of COMET offered by the ISPs. Content Consumers access freely to the content stored in content servers through the unified content access interface provided by the COMET system.

Figure 16 shows the entities involved in this business model and the payment relations. The blue arrows show payment flows, while also identify SLAs between the different entities.

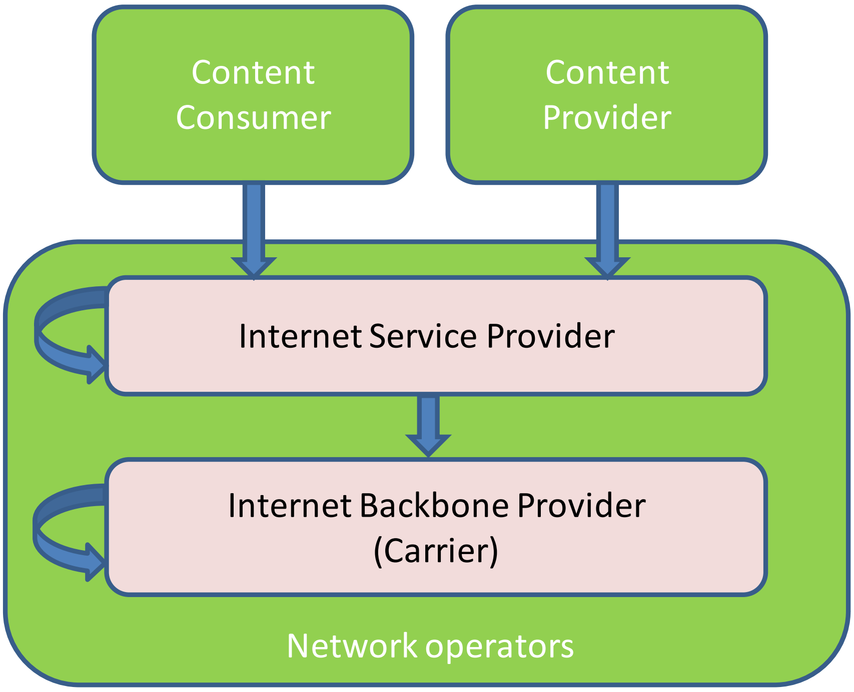

Figure 16: Entities and payment relations in the "Free content access" reference business model

The payment relations in this model are as follows:

- The *Content Consumer* pays the *COMET-aware ISP* for the ability (including content resolution and delivery) to consume contents from COMET, but it does not pay for the contents themselves, as they are free of charge.
- The *Content Provider* pays the COMET-aware ISP for content publication across the Internet and the content-aware distribution (anycast features, etc.). Again, the access to contents is free of charge.
- The *Network Operators* (ISPs and Carriers) will offer each other enhanced content delivery through COMET and will be paid accordingly (in a similar fashion as it happens today with data-based interconnection agreements). Besides, a different business model based on charging for COMET capabilities (e.g. content resolution, content register dissemination, etc.) could be followed.

5.2.2 Charged content access

Two reference business models of the COMET system can be identified where the Content Consumer is charged for content access, both requiring an enhanced connectivity provided by COMET-capable ISPs and carriers. The difference between both models lies in the recipient of the payments per content. Thus, in one option, the Content Consumer would also be charged by the ISPs (*ISP-based charged content access*), while, in the other model, the own Content Providers would be directly paid by the end user (*Content Provider-based charged content access*). Next, both models are described in more detail.

ISP-based charged content access

In this model, contents are published by Content Providers through the COMET publication capabilities, but they are offered to the end-users by ISP platforms and part of the global "COMET catalogue". Thus ISPs would charge Content Consumers for accessing the content, either on a pay-per-view model or on a subscription-based model.

Next figure shows the entities involved in this business model and the payment relations, where blue arrows show payment flows and SLAs.

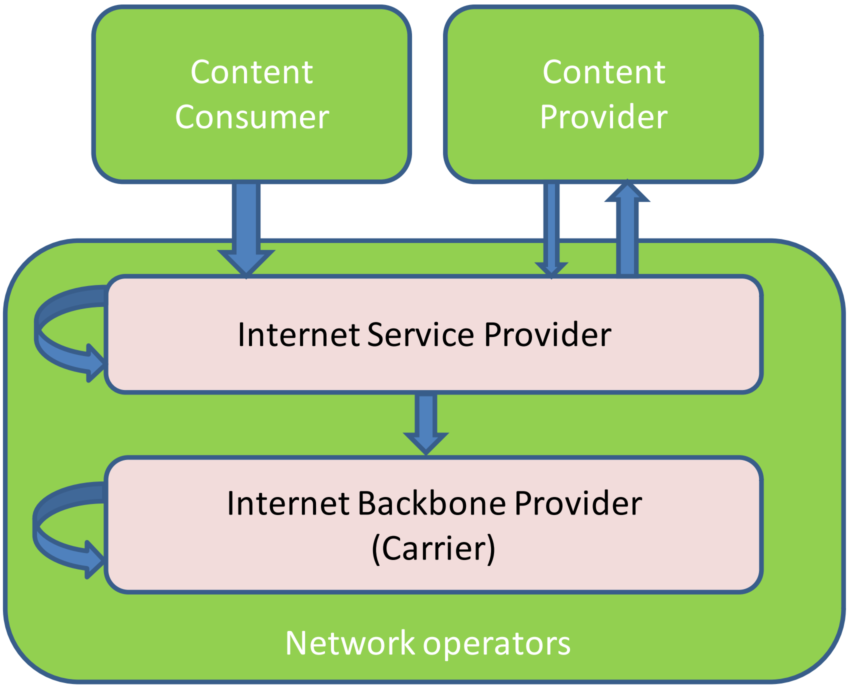

Figure 17: Entities and payment relations in the "ISP-based charged content access" reference business model

The payment relations in this model are as follows:

- The *Content Provider* pays the *ISP* for content publication (for making advertisements to attract potential Content Consumers)
- *Content Consumers* pay the ISPs for receiving the content. These payments can be included in the global bill for the COMET services offered by the ISP to the end user.
- The *ISP*, which collects all the payments from the end-user, pays back to the *Content Provider* for all the content accesses from users in its domain.
- As it happens with the Free Content Access, the *Network Operators* will offer each other enhanced content delivery through COMET and will be paid accordingly.

Content Provider-based charged content access

In this model, the content is published by *Content Providers* through the COMET publication capabilities and it is offered to the potential *Content Consumers* through *Content Provider*'s platforms or through *Content Distributor*'s platforms (*Content Distributors* were not included in the figure for the sake of simplicity). *Content Providers* can charge *Content Consumers* for the content directly, either on a pay-per-view model or on a (long-term) subscription-based model.

Figure 18 shows the entities involved in this business model and the payment relations, where blue arrows show payment flows and SLAs.

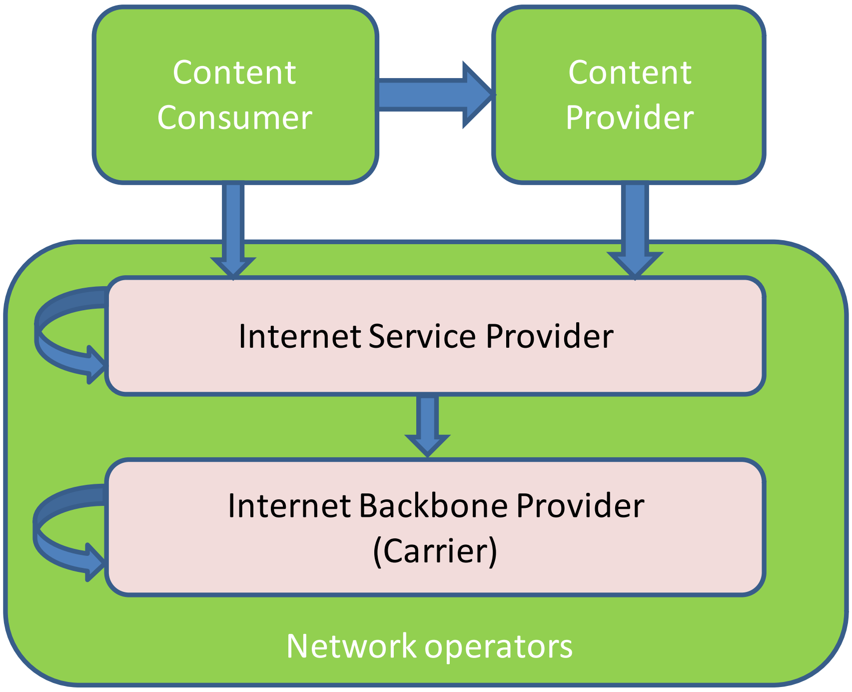

Figure 18: Entities and payment relations in the "Content Provider-based charged content access" reference business model

The payment relations in this model are as follows:

- The *Content Consumer* pays the *Content Provider* for the content (either on per-view basis or on long-term subscription bases).
- The *Content Provider* allocates part of the revenue from the *Content Consumers* and pays the *ISP* for content delivery, being possible for the ISPs to charge the Content Provider on a content basis (and not in terms of data traffic).
- The *Content Consumer* pays the *ISP* as well for providing the access to the COMET environment, independently from the content providers, but would not pay per content.
- As it happens with the other reference business models, the *Network Operators* will offer each other enhanced content delivery through COMET and will be paid accordingly. Besides, since Network Operators will be aware of every content request, business models based on cascading payments with charging on a content basis could be followed.

In some cases, such as in a long-term subscription to a commercial channel like Sky Sports, Content Consumers would probably make only one payment just to the Content Provider, including the content delivery in such payment. Content Providers will then allocate part of the revenue to the ISPs.

5.3 Business Models built over the Free Content Access Reference Model

These models are the natural evolution of the Over-the-top content distribution systems and their business models described in section 2.3, now enhanced by the COMET system. All these models fit into the free content access reference model.

5.3.1 Content Provider acting as Content Distributor

This section describes the evolution of those content distribution systems with providers acting as distributors from the current Over-the-top model to the more integrated model provided by the COMET system. Particularly, this business scenario has advantages such as the content delivery with guaranteed QoS and reduction of entry barriers for content publication.

Firstly, the effect that the COMET system has in the former business models is presented and, secondly, the new business models that could emerge from the implementation of the COMET system are introduced.

The COMET effect

Content Providers acting as Distributors are changing how consumers access to broadcast entertainment, and how content and communication companies provide it. Providers as YouTube or Dailymotion have become some of the major sources of traffic in the Internet. However, it is known that these Providers distribute their contents using best-effort techniques as they are usually not aware of the network topology they have below. What is more, they cannot provide any guarantee of delivery to the Content Creator itself as they do not have control over the underlying end to end transmission infrastructure. Some other disadvantages are the lack of control of the contents by the Creator and the requirements that the provider imposes to them to adapt their contents to specific format and characteristics.

COMET system will act as a mediator between these *Providers* and the *Network Operators* with their end-users. By doing this, it will help to solve the problems affecting the current platforms, which were introduced previously.

Content Consumers, when acting as Providers (*prosumers*), will benefit from the COMET mediation as well, since they will have a unified way of publishing their contents, solving the problem of current Internet, where each intermediary or each *Content Provider* offers a different way of publication to the *Content Creator*. This unified interface with the COMET system will also offer to *Content Creators* or *Providers* the capability of controlling their contents (e.g. popularity measurements, content characteristics updating…). In addition, the *Consumers* will experience a better grade of service while the *Providers* are benefiting from the COMET mediation and will be able to offer an enhanced service.

COMET system can also offer a better distribution system to Content Providers that have their own way of reaching the end-users. By the capabilities offered by the CMP to these Providers, benefiting from its awareness of the network topology (e.g. finding the best path from the server to the end-user), these platforms could achieve a more efficient distribution, being able to guarantee a given grade of service to their customers.

Network Operators can benefit from the mediation that the COMET system introduces as this mediation will be aware of the network topology and can take into account Operators' economical and business policies.

Summarizing, the COMET system will help to solve most of the problems that current Content Providers and Content Creators currently have to distribute their contents through the Internet efficiently.

Evolution of business models

Figure 19 depicts the business relations that are present in the systems in which *Content Providers* offer and distribute free contents to the *Content Consumers*. The figure on the left represents the roles and the figure on the right introduces an example of each kind of player. Additionally it is shown a "*prosumer*" player, which represents an amateur *Content Creator* who makes use of the services of the *Content Provider* and *Distributor* to publish their contents. The relation it establishes with the Content Provider is free in this business model. For that reason, there is no cash flow arrow between these roles in the figure.

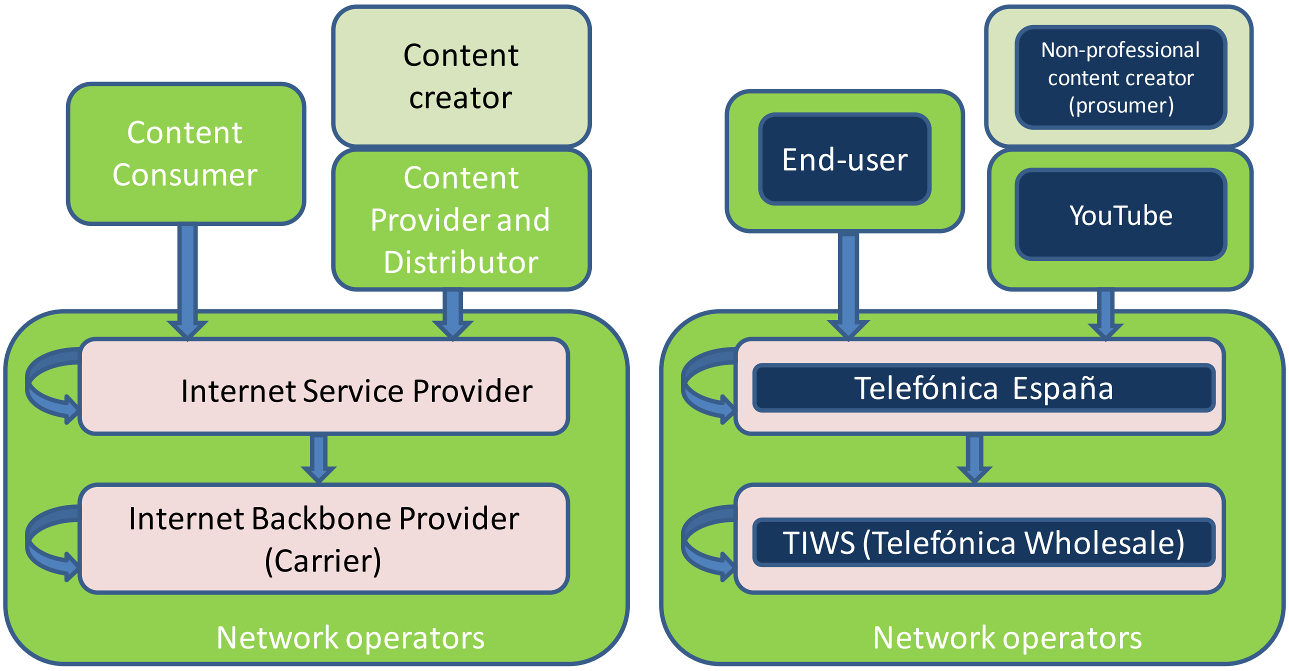

Figure 19: Entities and payment relations in the content distribution from a merged content provider and distributor

The existence of the COMET system could change the current business models of the *Content Providers* which are also acting as distributors. With the functionalities that COMET provides to these Providers, they will be able to offer an enhanced service to their customers so they could charge them for it or create a differentiation between standard and premium consumers.

The business model that involves the relation between *Content Providers* and *Network Operators* could change in both ways. On the one hand, if the Network Operator providing the Internet access service to the *Content Provider* were also part of the COMET system, it could offer an enhanced connectivity service, with facilities of publication, traffic reduction (e.g. due to multicast techniques, etc.), so that these advanced functionalities could be included as additional payments. On the other hand, the *Content Providers* which decided to hire the COMET mediation service might benefit from a traffic reduction in their Internet access (e.g. not all the content replicas should be created by the source, thanks to multicast) so that they could reduce the part of the bill related to the amount of bandwidth served in that Internet access.

Content Providers acting as *Distributors* will definitely benefit from the deployment of the COMET system, as they would be able to provide higher quality and more reliable content delivery to their customers, even with the possibility of an increase of revenue from them in the premium segment.

5.3.2 Content Delivery Networks

This section describes the evolution of the Content Delivery Networks from the current Over-the-top model to the more integrated model provided by the COMET system. As it will be shown, this scenario presents advantages from the integration of CDNs with the underlying network operators' infrastructure, such as an improvement in the consumer's perceived QoS.

In this section, the effect that the COMET system has in the former business models is presented as well as the new business models that could emerge from their integration into the COMET system.

The COMET effect

In recent years, the emergence of CDNs has brought a convenient solution to enhance the quality of the Internet. As it is widely known, a CDN can be understood as a collection of network equipment offering the same content (by replicating it to multiple locations) to the end-users. Previously, the original content has to be replicated and located on a server, called the Origin Server. The replication of the content is based on various techniques of content caching. One of the key purposes of the CDN mechanisms is to be able to cache the content on different servers in an efficient manner and route the requests from the end-users to the most adequate of the available servers. It should be noted that one of the advantages of this architecture is to reduce traffic to and from the origin server, since massive contents requests can be served directly from the network where the end-user is subscribed.

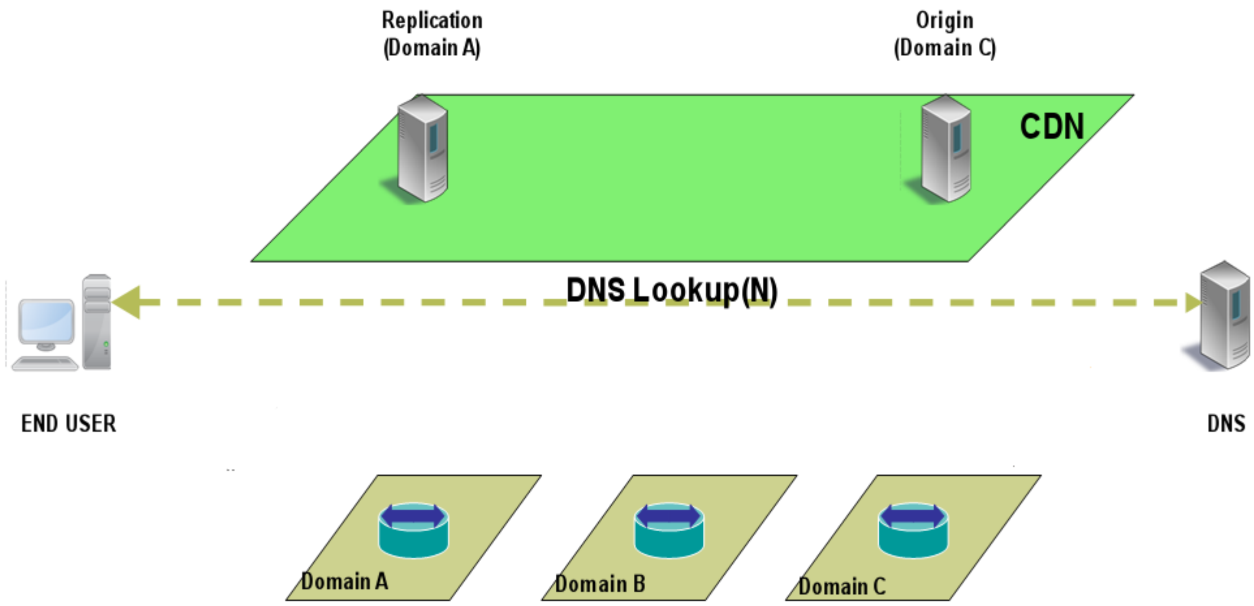

Figure 20: Current CDN architecture

Figure 20 shows the loose relation between the current CDNs and the transmission infrastructure provided by Network Operators. As it is depicted in the figure, there is no interface between the CDN and the underlying network. Currently, due to the fact that this interface is missing, the decisions in the CDN plane cannot be coordinated with the decisions in the underlying network infrastructure, and vice-versa. Hence, a major drawback of this architecture is the fact that conventional CDNs can only support best effort services, as no effective coordination can be established with the QoS mechanisms currently available in the network [1].

The COMET project studies and investigates the disadvantages mentioned above, in order to find solutions based on a more suitable naming architecture and improved routing and traffic engineering functionalities natively connected to the content retrieval procedure. Thus the two planes of the COMET system (*Content Mediation Plane* and *Content Forwarding Plane*) will be able to enhance the DNS-based resolution of the CDNs while orchestrating a coordinated delivery strategy with the underlying network infrastructure.

In order to achieve that coordination from the beginning, the CDNs would be able to inform the COMET's CMP about the content and the location of its replicas during the content publication phase. According to this process, the CMP would be aware of the content location and could effectively use performance information from the network layer to adapt the routing according to content availability and quality of service requirements, in coordination with COMET's CFP.

In summary, the COMET architecture, through the two planes, will be able to provide the missing link between the two different "worlds" —"contents" vs. "data transmission"— represented today by the CDNs and the Network Operators. Using these COMET planes, both "worlds" can share information efficiently (on content availability, routing, QoS, etc.) to be used in cooperation with COMET algorithms in favour of *content providers* and *consumers*.

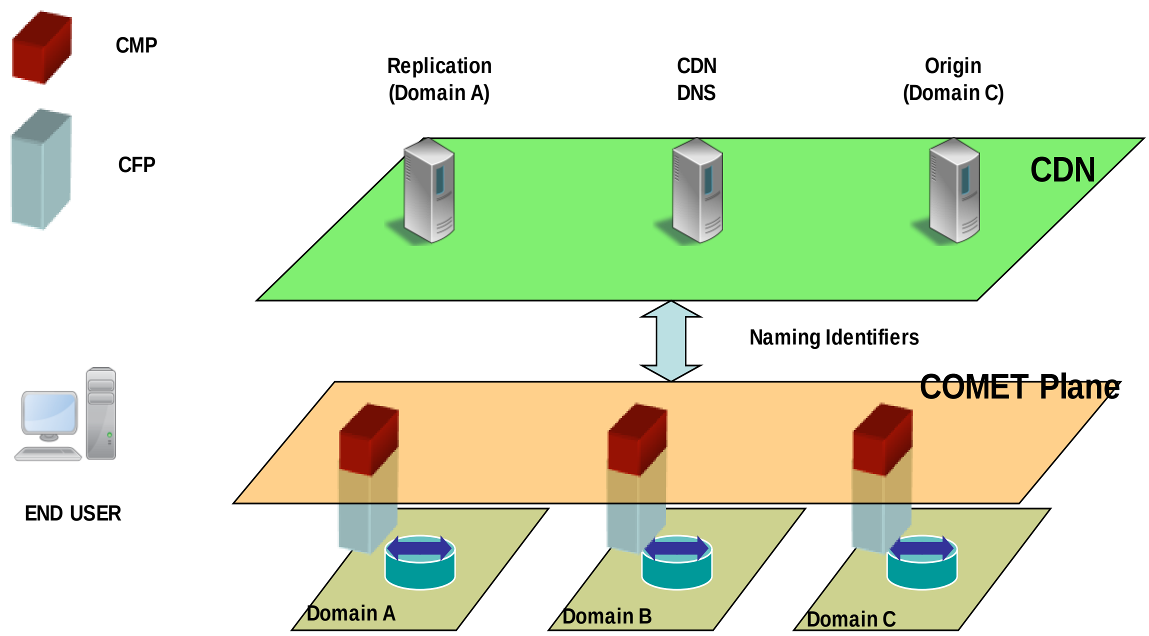

Figure 21: The COMET system in CDN Architecture

Evolution of business models

Currently there are many companies wishing to participate in the CDN market, commercialising their own branded CDN service, but cannot afford the high entry barriers associated to the development and deployment from scratch of their own proprietary CDN architecture and the high transaction costs associated to gain a minimum amount of contents. However, being the COMET architecture in place, these entry barriers for new agents operating in the replication and hosting arena could be highly reduced, as a standard Internet system such as COMET might be in charge of all the logic related to content publication, content retrieval, worldwide content access and even the QoS assurance for the transmission (which is not possible in the current OTT approach). Thus, to become a new CDN player, acting as a big aggregator of contents, the minimum technical requirements would just be the maintenance of a set of hosting servers efficient and the efficient replication of the contents among them.

On the other hand, incumbent CDN players could also evolve to COMET-aware architectures where they could benefit from the traffic engineering techniques available in COMET's CFP while being fully integrated in a unified content naming scheme.

In both cases, the business model participated by these COMET-aware CDN players would consist of 4 entities: the *Content Creator* —responsible of producing the content (broadcasters, studios etc.)— the *Content Consumer* —that is the end-user—, the COMET-aware CDN provider, and, finally, the ISP. In this model, the *Content Provider* (i.e. studios) would pay the CDN for hosting, replicating and distributing their contents. The CDN, in turn, would pay *Network Operators* for providing a wholesale Internet access with COMET capabilities, while *Content Consumers* are also paying their local ISPs for their COMET-enabled Internet access. Finally, operators might settle interconnection agreements with intermediate carriers, paying in turn for that service. Figure 22 shows a general overview of the entities and payment relations involved in these business models. The part on the left represents the generic business entities while the part on the right introduces an example of players.

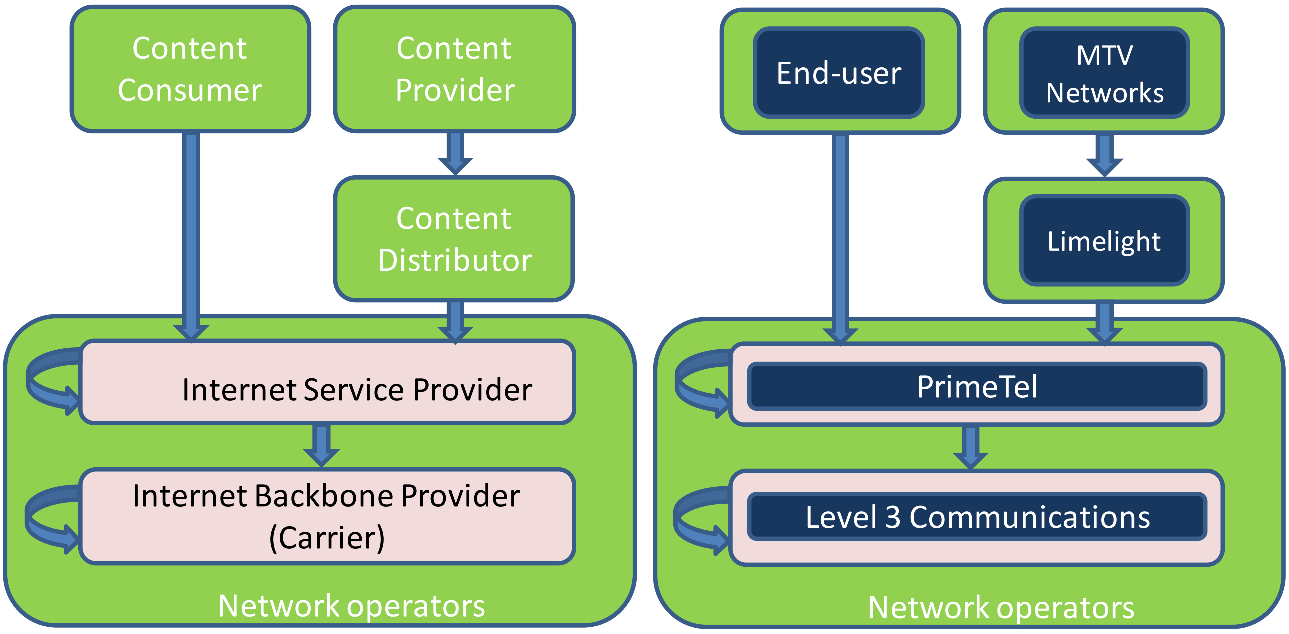

Figure 22: Entities and payment relations in the content distribution through CDNs with COMET

5.3.3 P2P Networks

This section describes the evolution of the P2P Networks from the current Over-the-top model to the more integrated model provided by the COMET system. Particularly, this business scenario has advantages as the removal of barriers in content search as well as a significant cost reduction for several Internet actors.

Firstly, it is discussed the effect that the COMET system would have over the previous business models and, next, the new business models that could emerge from the deployment of the COMET system are presented.

The COMET effect

P2P networks have established their position in Internet during the last decade, by consuming a vast amount of Internet traffic and transforming Internet business models accordingly. Generally, P2P networks have specific disadvantages such as the lack of guarantee for QoS, or the inefficient directory lookup in some cases, which are being investigated and might be addressed with their inclusion in the COMET framework.

The COMET system will act as a mediator between P2P platforms and the rest of the network. Particularly, it will allow the integration of the wealth of contents currently available in P2P networks into a unified catalogue of contents accessible by COMET's *Content Consumers* via the unified COMET interface.

Either causal "prosumers" or other kinds of *Content Providers* will be able to publish their contents in COMET and distribute them via P2P platforms, by allowing the content to be stored and re-distributed. In consequence, content search can be simplified and unified by using the COMET system, and P2P networks would only be responsible for the distribution of the content.

Moreover, content publication through the COMET system could also lead to a more reliable content identification in P2P networks, as the *Content Consumers* would receive the content they had requested (and not a fake one), which brings significant advantages. Registration of content in the COMET system, along with the registration of each *Content Provider* himself, would lead to a more secure distribution of contents in the Internet, providing more reliability to P2P platforms.

Apart from the advantages described above, the COMET system, and specifically the CMP, might be aware of P2P networks and peers participating in swarms. In that case, the COMET system might even have the ability to monitor their location and status, along with information about the requesting user, and possibly help the end-user application to decide the best sources (e.g. preferring peers in the same network rather than peers in external networks). However, this functionality, although theoretically compatible with the COMET approach, is out of the scope of the project and might be a matter of further study.

In conclusion, the COMET system could solve significant disadvantages of P2P networks through the new capabilities for content publication and search, targeting to provide a better and more homogeneous content reachability for end-users, even improving content reliability.

Evolution of business models

Figure 23 depicts the business relations between the entities that are present in the scenarios where *Content Providers* use P2P platforms to distribute free contents to *Content Consumers* via the COMET system. The part on the left represents the generic business entities while the part on the right introduces an example of players.

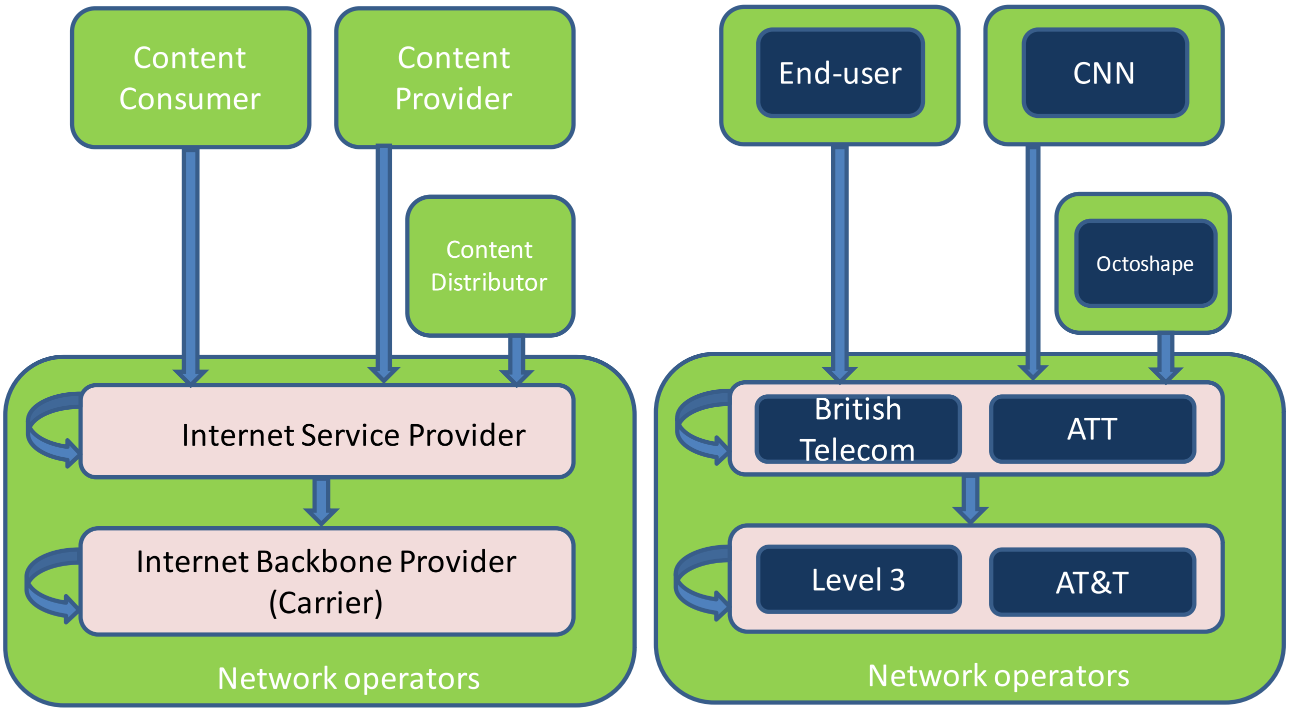

Figure 23: Entities and payment relations in the content distribution through P2P networks with COMET

The entities and payment relations in this scenario are similar to the ones presented in the free content access reference model (section 5.2.1), with the exception of the *Content Distributor* (P2P Platforms). As presented in the reference model, *Network Operators* adopting the COMET system could charge their customers with higher fees for an enhanced access to content through the unified COMET interface, which now includes the access to contents distributed via P2P networks.

P2P content providers will definitely gain noticeable benefits from the adoption of the COMET system, as their contents will be more reliable and more easily reachable and searchable in the Internet.

5.4 Business Models built over the Charged Content Access Reference Model

This chapter presents the business models where the *Content Consumer* is charged for retrieving contents. These models fit into the charged content access reference models described in section 5.2.2.

Nowadays, with the current Internet architecture it is not possible to distribute contents with the end-to-end quality achieved by traditional broadcasting platforms (DVB-T, DVB-S, etc.). This architecture has technical problems that prevent an Internet user to watch high quality live contents smoothly (e.g. HDTV through the Internet). In addition, the current solutions that are close to achieve this required QoS, incur substantial costs for being able to offer it to wide audiences.

For these reasons, in the current Internet there is a lack of business models based on payments for the live broadcasting of contents. The existing business scenarios are restricted to the following ones:

- Live (non high quality content) broadcasting for free.
- Non-free downloads of high quality content to the Content Client (e.g. iTunes Store, Pixbox) for later viewing.

In the cases in which not much bandwidth is needed (e.g, Spotify), a high-quality live contents distribution is possible but, when there are stricter requirements of bandwidth for the distribution and this bandwidth must be shared between various end-users, smart mechanisms must exist to guarantee a particular level of bandwidth per user in order to provide a smooth viewing.

Given that the ISPs will become integrated in the COMET system architecture, the involved players could take benefit from the following advantages:

- As the Network Operators are aware of (and control) the network topology, and given the interconnection agreements that they may have between each other, they will be able to provide end-to-end QoS to the Content Providers.
- The Content Providers will require less access bandwidth to simultaneously distribute their contents to wide audiences due to the multicast techniques that Operators could implement in their networks.
- The publication process will be simplified for Content Creators and Content Providers. With the deployment of the COMET system it will be easier for them to publish their contents through the Internet.

The COMET system will help the emergence of an Internet model in which Network Operators will globally offer similar services to those offered by current CDNs, but in a generalized way to all users, so that:

- Content Providers are able to offer their contents via the Internet and with low bandwidth requirements.
- Content Creators are able to become Content Providers with minimum infrastructure deploy costs.

In this way, there are two possible business scenarios with payments associated to the content viewing: Internet broadcast services, in which the Content Provider directly charges the Content Consumer, and IPTV managed services, in which the Network Operator agglutinates the payments from the Content Consumer and allocates part of this payment to the Content Provider.

The business models and payment relations between the entities of these scenarios are detailed in the following sections.

5.4.1 Internet TV broadcasting services

Internet TV broadcasting industry has started to take off as it offers more services to its customers with the main change being Digital television and interactive applications. This revolution in the broadcasting market adds extra load to the network, as it requires both a high capacity and good Quality of Service (QoS). It is apparent that there is a need for an improvement in the way that contents are delivered. COMET comes to give the solution to the current content delivery limitations. With COMET redefining content delivery, Internet based broadcasting services, such as VoD or Live streaming services, would benefit from an improved delivery both in terms of flexibility and performance.

Figure 24 depicts the business relations between the entities that are present in the scenarios that offer Internet broadcasting services. These business models and the payment relations between the entities will be described further below. The figure on the left represents the roles, and the figure on the right introduces an example of players.

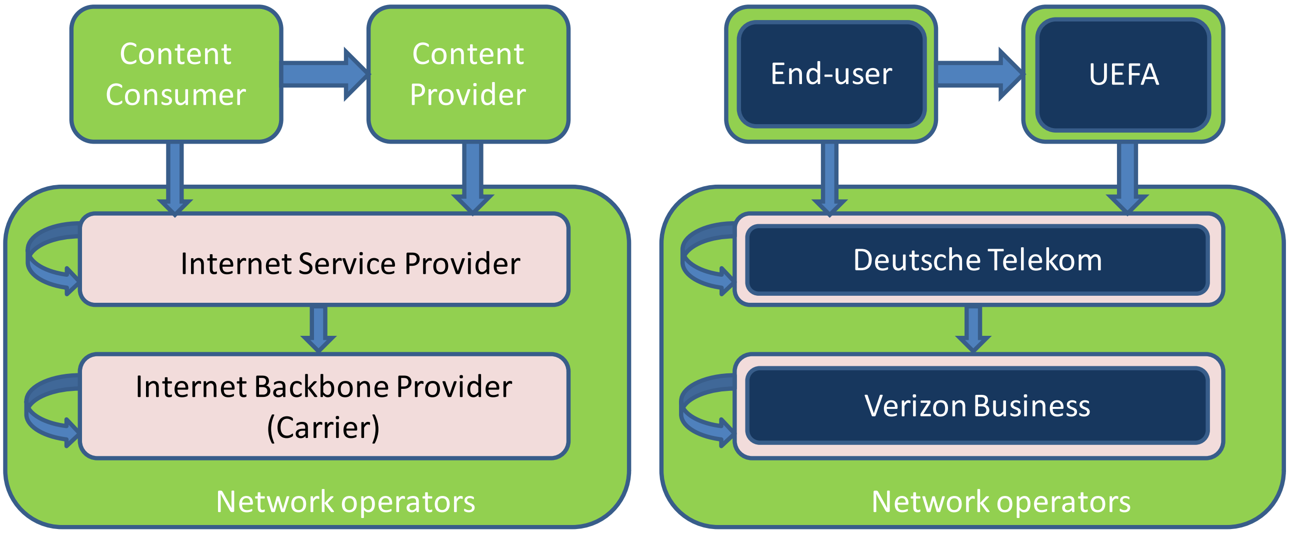

Figure 24: Entities and payment relations in the content distribution through Internet Broadcast services with COMET

The relations that are shown in the figure are the following:

- *Content Consumers* pay *Content Providers* for accessing their contents. This relation could be agreed on a long-term payment (e.g. monthly) or based on a per-content payment (e.g. pay-per-view model).
- *Content Providers* allocate part of the revenue from the *Content Consumers* and pay *ISPs* for the content delivery. In this scenario, the *Network Operator* could act as *Content Distributor* for the *Content Provider*.
- *Content Consumers* pay the *ISPs* for providing the access to the COMET environment, independently from the *Content Providers*, but would not pay per content.
- The *Network Operators* will offer each other enhanced content delivery through COMET and will be paid accordingly.

5.4.2 IPTV managed services

The *ISPs* must nowadays face the competition of Over-the-top (OTT) platforms in the contents distribution market through the Internet. In the last years, these Network Operators have developed their own IPTV systems in order to participate in the profits that are present in this market.

However, given the dynamism and ease of OTT platforms to offer and attract contents from the Creators and Providers, ISPs are losing ground with respect to them. Specifically, these IPTV platforms have a reduced catalogue of channels, usually composed by the most popular broadcasting companies, but these are not enough to accurately meet the end-users demand. In that way, the minority contents are not accessible through these platforms, so the long tail present in this market is not exploited and a large share of the population are not able to satisfy their content requests.

Another disadvantage for IPTV platforms is that they must make a proactive search of contents reaching agreements with the different Content Providers and Creators in a *pull-based* model. After paying for all these contents, the ISPs act as aggregators and offer the contents to their clients.

Thanks to the ease of content publication through the COMET system, *ISPs* would have the ability to deploy a more open IPTV model (OpenIPTV) in which the niche contents could be offered and all Content Creators and Providers would be able to distribute their contents through this new platform. This model follows the trend started by some ISPs of exporting their own capabilities to third-party agents (e.g. OpenTelefonica, BT Web21C API) only that in this scenario it is the distribution channel what is opened.

The relation between the Content Provider or Creator and the IPTV platform would involve an agreement to show the published content. The ISP's IPTV platforms will have to pay for the rights of the stream (similar to what happens today) and will have the right to use a unique content name to show the content to its clients and charge them accordingly.

The advantages from the scenario of the IPTV enhanced with the COMET system are the following:

- ISPs would offer a larger and more attractive catalogue of channels through their IPTV platforms. This catalogue would meet better the demands from their clients by offering them the contents that they want to request.
- The ISPs could focus their attention in reaching agreements only with the Content Providers that offer the contents that their clients are demanding.
- The adoption of the COMET system favors a revenue sharing model in which the Content Consumers pay for the viewing of the contents.
- Small or minority Content Creators could offer their contents through the IPTV platform enhanced with the COMET system and take revenue from the provision of these contents.

In addition to these advantages, by the adoption of the COMET system, the ISPs, with their IPTV platforms, could be able to offer contents to Content Consumers from outside their networks, that is, not being their clients. The COMET system will provide interconnection capabilities that could make possible inter-AS content distribution.

Figure 25 depicts the business relations between the entities that are present in the scenarios where the content is offered by IPTV managed services to the Content Consumers. These business models and the payment relations between the entities are described further below. The figure on the left represents the roles while the figure on the right introduces an example of players.

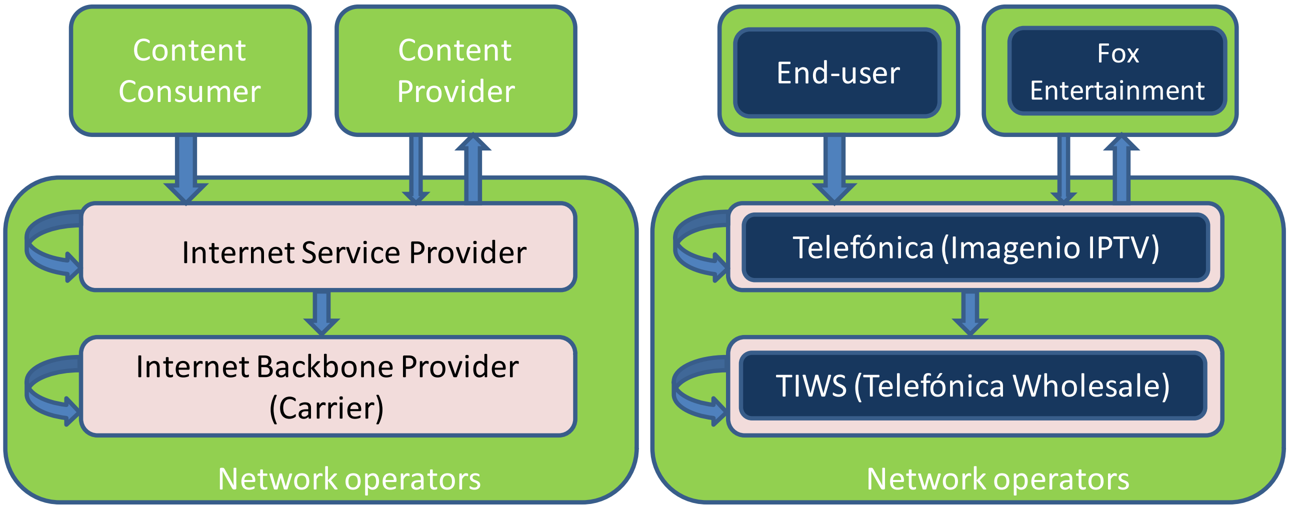

Figure 25: Entities and payment relations in the content distribution through IPTV managed services with COMET

Payment relations shown in the figure are the same as those that were described in the ISP-based charged content access (section 5.2.2):

- *Content Consumers* pay the *ISPs* for accessing their contents and for providing the access to the COMET environment, independently from the *Content Providers*. This relation could be agreed on a long-term payment (e.g. monthly) or based on a per-content payment (e.g. pay-per-view model).
- *ISPs* allocate part of the revenue from the *Content Consumers* and pay the *Content Providers* in a revenue sharing model.
- The *Network Operators* will offer each other enhanced content delivery through COMET and will be paid accordingly.

5.5 Impact of COMET in data interconnection business models

This section is intended to show the impact that the COMET system can have in QoS and multicast interconnection business models. It is not the aim of the COMET project to define new business models related to Layer-3 interactions, but to offer a flexible umbrella where different QoS solutions (AGAVE Network Planes and Parallel Internets [24], Diffserv-based networks [25], etc.) can fit, and where point-to-multipoint content delivery is possible.

In this way, we expect the COMET system to be a driver for Network Operators to establish QoS and multicast interconnection agreements and boost the development of a QoS- and multicast-aware Internet.

5.5.1 Impact of COMET in QoS interconnection

Technical issues on QoS interconnection

While Quality of Service (QoS) provisioning techniques within a single autonomous domain have been relatively mature, how to achieve end-to-end QoS across multiple ISP networks is still yet to be obtained. The major difficulties in enabling Internet-wide QoS basically lie in the following two aspects:

(1) Scalability: given the existence of more than 30,000 autonomous domains today, how they can establish QoS-enabled pipes in a scalable way remains unclear. The Path Computation Element (PCE) paradigm has offered a promising platform that supports computing explicit end-to-end Label Switched Paths (LSPs) across multiple domains, possibly with QoS awareness, how such a path computation “facility” can be driven by tangible QoS interconnections between domains is still largely unknown.

(2) ISP policy heterogeneity: it is well known that different ISPs have their specific policies in managing own network resources. For instance it is difficult to conceive how end-to-end QoS can be achieved through two adjacent domains, with one of them focusing on minimising edge-to-edge delay, while the other provisions its resources for maximising throughput. In this case, some type of negotiations between autonomous ISPs is needed if they want to participate in enabling end-to-end QoS with consistent goals.

Current QoS interconnection business scenarios

Generally, it is required that ISPs negotiate provider-level service level specifications (pSLSes) with each other for QoS interconnection in the Internet. The establishment of pSLSes is normally driven by the business objective (e.g. service level targets) and network resource provisioning policies set by individual ISPs. As far as QoS interconnection models are concerned, there exist two popular scenarios proposed in the literature: the centralised model and the cascaded model ([9][10]).

Centralised Model

This model is also known as source-based model. The main idea is as follows. For any source domain that needs to establish an end-to-end QoS pipe to another remote domain, it needs to contact all the intermediate domains towards that destination domain, and establish a dedicated pSLS with each of them. Figure 26 illustrates a simple example for the idea. Each autonomous domain (AS1 – AS3) locally engineers its local edge-to-edge QoS capabilities, represented by l-QC1, l-QC2 and l-QC3 respectively. If AS1 wants to establish an end-to-end QoS pipe to the remote destination AS3, it needs to negotiate with all the intermediate ASs (in this case AS2 only) as well as the destination AS. In Figure 26, the corresponding provider level SLS include pSLS1 with AS2 and pSLS2 with AS3 respectively. As the result, the end-to-end QoS capability (e-QC1) is established from AS1 to AS3, with the "concatenated" local QoS capabilities from individual domains along the pipe.

Despite the straightforward idea, it can be easily inferred that such an approach fundamentally suffers from scalability problem, and hence is not appropriate for universal end-to-end QoS across a large number of participating ISP networks.

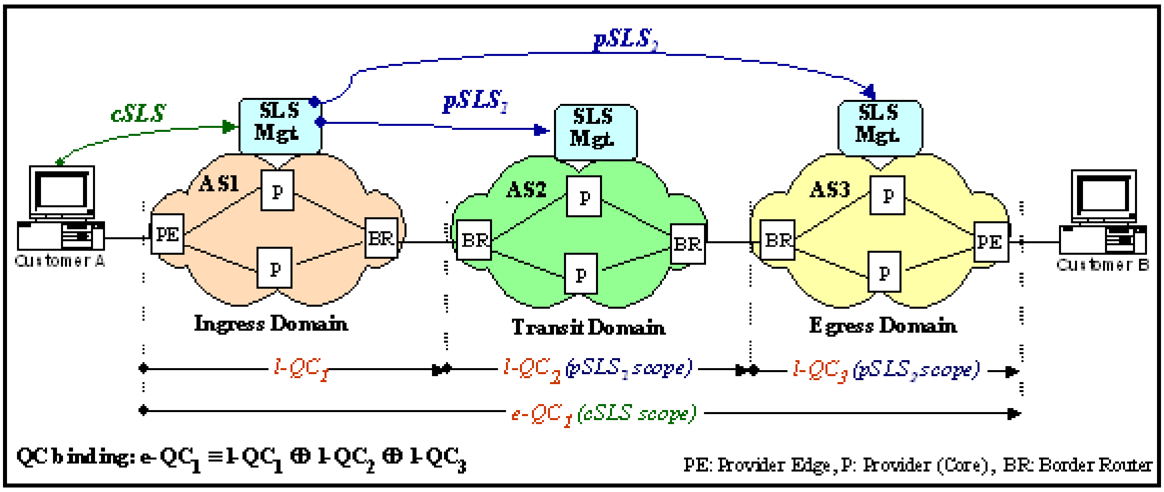

Figure 26: The centralised QoS interconnection model

Cascaded Model:

A much more scalable QoS interconnection model has also been proposed in the literature, commonly known as cascaded model. According to this model, each autonomous domain establishes pSLS only with the immediately adjacent ISP networks. As such, the QoS peering agreements are between adjacent neighbours and not between remote providers more than one hop away.

Figure 27 presents a simple example of such an approach. The local QoS capabilities engineered by ASs 1, 2 and 3 are denoted with l-QC1, l-QC2 and l-QC3 respectively. AS2 can negotiate a contract (pSLS2) with AS3, enabling AS2's local customers to reach destinations in AS3 with an extended QoS capability e-QC2 which is "concatenated" by l-QC2 and l-QC3. Thereafter, AS2 may advertise this e-QC2 capability towards destination AS3 to its own neighbouring domain AS1. Another provider level contract can be established between AS1 and AS 2, which is pSLS2. pSLS2 basically allows AS1's local customers to reach AS3 destinations with QoS awareness, and from Figure 27 we can see that the end-to-end QoS capability from AS1 to AS3 (denoted by e-QC1) is the concatenated QoS capability of l-QC1 and e-QC2.

The advantage of the cascaded model is scalability in the sense that QoS negotiations are on hop-by-hop bases and it is not necessary for the source domain to directly interact any remote ISP networks in order to establish inter-domain QoS capability.

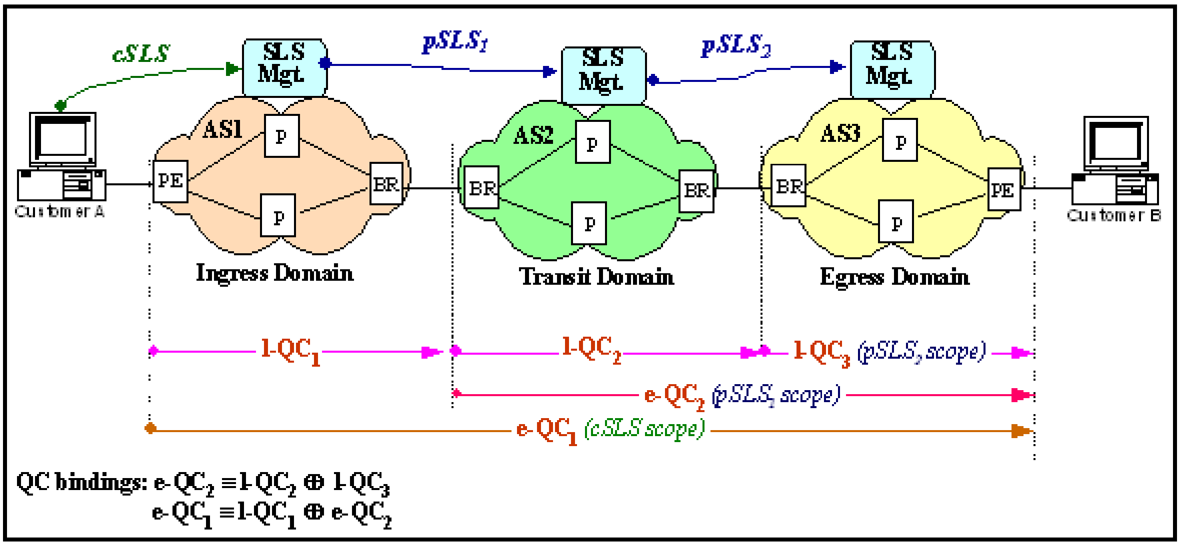

Figure 27: The cascaded QoS interconnection model

Impact of COMET

In spite of the numerous research work and effort on developing QoS interconnection business models [10][26][27][28] and end-to-end QoS technical solutions [10][28], these have only reached a reduced success in the Internet, limited to specific agreements between two interested parties, and never being Internet widespread. In fact, nowadays there is not a QoS-aware Internet able to support end-to-end content delivery with QoS guarantees. Network Operators have developed their own inter-domain QoS solutions such as the IPX model [29], but their use is restricted to operator services such as VoIP and videoconference and is not open to Internet services.

The capabilities of the COMET system to mediate content requests in the content mediation plane and the existence of Content-Aware Forwarders, specific network elements in the content forwarding plane able to deal natively with content packets, allow the application of QoS policies per content so that it is possible to offer guarantees to content delivery. It is not the aim of the COMET project to re-invent or design new end-to-end QoS strategies, but to re-use existing QoS solutions (AGAVE Network Planes and Parallel Internets [10], Diffserv-based networks [25], etc.). Specifically, the COMET system will rely on QoS solutions that advocate for cascading business scenarios where QoS aware content delivery can be offered in a scalable way.

5.5.2 Impact of COMET in multicast interconnection

Issues on deploying Internet-wide multicast

In the last twenty years many efforts have been done in order to deploy across the Internet multicast capabilities due to the advantages this technology provides:

- Network optimization. With multicast techniques it is not necessary to send a copy of the content per destination, since one data flow can be shared between multiple destinations, avoiding unnecessary data replication and, therefore, obtaining bandwidth savings. Moreover, a linear increase in the number of requesters/destinations does not produce a linear growth of network traffic. In this way, networks can deal better with crowded real-time services and applications.
- Optimal resources performance: With multicast, the number of copies of a content sent from a content server is limited to only one. This reduces servers CPU load and bandwidth, thus allowing serving more contents and destinations per CPU.

Thanks to these advantages, multicast turns out to be much more suitable than unicast for broadcast applications and services such as Web-Inars and Internet TV. For this reason, multicast is a recurrent research topic and there is a growing interest on making it available across the whole Internet.

However, in spite of the previous advantages, the current Internet is not able to bear efficiently worldwide multicast services or applications. Technical reasons lie in the fact that, in order to deploy an inter-domain multicast network, two requirements must be met:

- All routers in the path between multicast sources and end users subscribed to multicast groups are required to be IP multicast forwarders.
- All domains which have multicast sources or potential multicast end users must participate in the multicast routing. Regarding this issue, some working groups such as MBone [19] or Internet2 [20] have tried to create an overlay multicast network through unicast IP tunnels configured between multicast-native domains, but their success has been restricted to experimental environments and has not been extended to the whole Internet.

Besides these technical inconveniences, there are also business issues which have possibly influenced the adoption of this technology at a worldwide scale:

- The current business models related to multicast traffic do not take into account the induced costs per domain produced by multicast data distribution. These models should consider the fact that one copy of a content in the source domain can generate multiple copies of this content in intermediate and destination domains. Specifically, the traffic in the end users' access links remains the same, so destination domains incur in most of the costs for delivering the content to the end users.
- A Network Operator involved into a multicast data transmission is not able to check the number of content replicas generated in the next neighbour Network Operators. This prevents that Network Operators can charge multicast traffic depending on the number of content replicas and, therefore, makes impossible the development of business models which capture the actual costs of delivering multicast traffic.
- Finally, carriers offering wholesale data transport services do not have incentives to implement multicast capabilities within their own network, because of the decrease of revenues that would come from a traffic reduction from their customer-ISPs. Next, a more detailed explanation can be found.

Lack of incentives for carriers to deal with multicast traffic

Figure 28 represents a general model of the current unicast data interconnection business model and the payment relations between Network Operators.

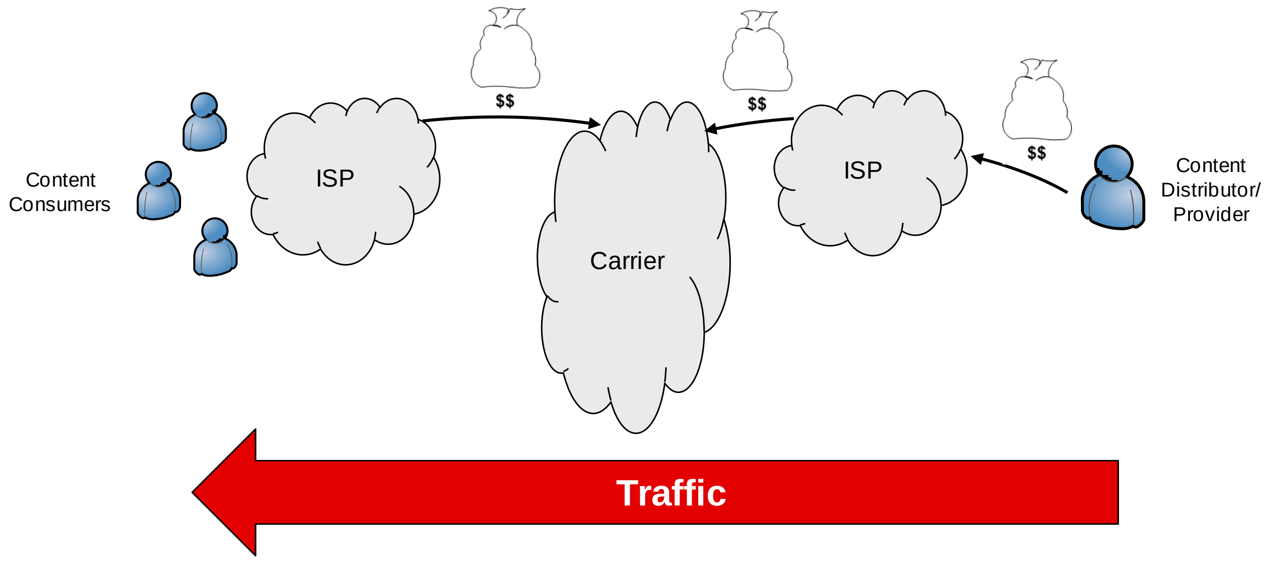

Figure 28: General interconnection business model (unicast)

Figure 28 shows just a common scenario of payment relations between actors. Other payment relations are possible, although they are not shown in Figure 28 due to the heterogeneity of interconnection agreements (for instance, ISPs usually have peering relationships with other ISPs where they do not pay each other).

As it can be seen, Content Distributors/providers have to pay to a local ISP for providing them access to the Global Internet, depending on the traffic volume they generate towards the Internet. On the other side, ISPs are charged by carriers for transporting their data traffic from and towards the rest of the Internet.

If the current unicast interconnection business model were applied to multicast traffic, a traffic reduction would be expected and the payments would change according to Figure 29. Green dialog boxes show the benefits for the different actors while red ones emphasize actors' detriments.

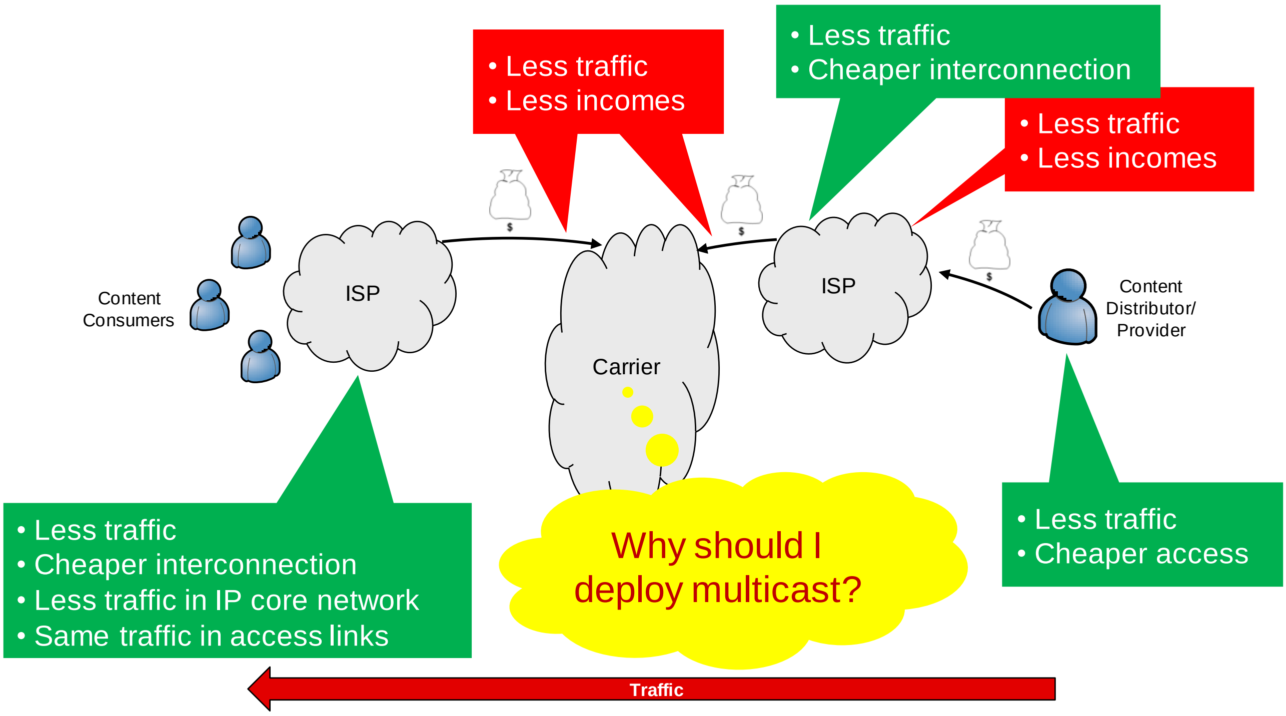

Figure 29: Impact of multicast deployment in current interconnection business model

Content Distributors/providers achieve a reduction in terms of server load and traffic generation, obtaining remarkably bandwidth savings. These bandwidth savings allow them to reduce the capacity of its Internet access, achieving cost savings too.

Content Consumers' ISPs achieve costs savings due to the reduction of traffic from the rest of the Internet and also within their own IP core networks. It must be noticed that multicast traffic does not imply a reduction in the end users' access links. In the same way, Content Providers' ISPs achieve costs savings due to the reduction of traffic towards the rest of the Internet, although they will also receive less traffic from Content Providers and therefore, will obtain lower incomes.

Finally, Carriers do receive less traffic both from Content Consumers' ISPs and Content Providers' ISPs, decreasing dramatically their revenues. For this reason, Carriers do not have incentives to support multicast within their networks.

Business models related to point-to-multipoint content delivery with COMET

As stated in the DoW, it is not the aim of the COMET project to re-invent the multicast technology, but to make use of the existing multicast technologies and provide an integrated solution where the COMET system can support point-to-multipoint content delivery thus achieving the advantages it offers (e.g. bandwidth savings due to traffic reduction and optimal resource utilization).

In particular, thanks to the capabilities provided by the COMET system, ISPs can be aware of the content requests, and, specifically, aware of those requests related to real-time content, more suitable to be distributed with multicast technologies. COMET ISPs can configure the Content Forwarding Plane (CFP) in an efficient way in order to deliver the content from the content server (one point) to the end users interested in that content (multiple points). Moreover, due to the content awareness, all ISPs involved in a point-to-multipoint transmission could be aware of the active content flows in the next neighbour ISPs. Therefore, it becomes possible to develop business models for point-to-multipoint content delivery where payments can be based on incurred costs and where charging can be verified by all the interested parties.

In this way, we expect the COMET system to be a driver for Network Operators to establish point-to-multipoint interconnection agreements and boost the development of a "multicast-aware" Internet. In this respect, two different point-to-multipoint interconnection business scenarios are foreseen:

- Carriers not supporting point-to-multipoint content delivery
- COMET-aware Carriers supporting point-to-multipoint content delivery

<u>Scenario 1: Carriers not supporting point-to-multipoint content delivery</u>

In this scenario, Carriers do not implement capabilities to support point-to-multipoint content delivery within their own networks. This model responds to the lack of incentives from Carriers to deploy multicast. Although Carriers do not support point-to-multipoint content delivery, the COMET system can still support it between edge COMET ISPs. Figure 30 illustrates this scenario, where Content Providers' ISPs would reach interconnection agreements with Content Consumers' ISPs to deliver only one content for all Content Consumers in that ISP. For the Carriers, the COMET traffic will be considered just as unicast traffic

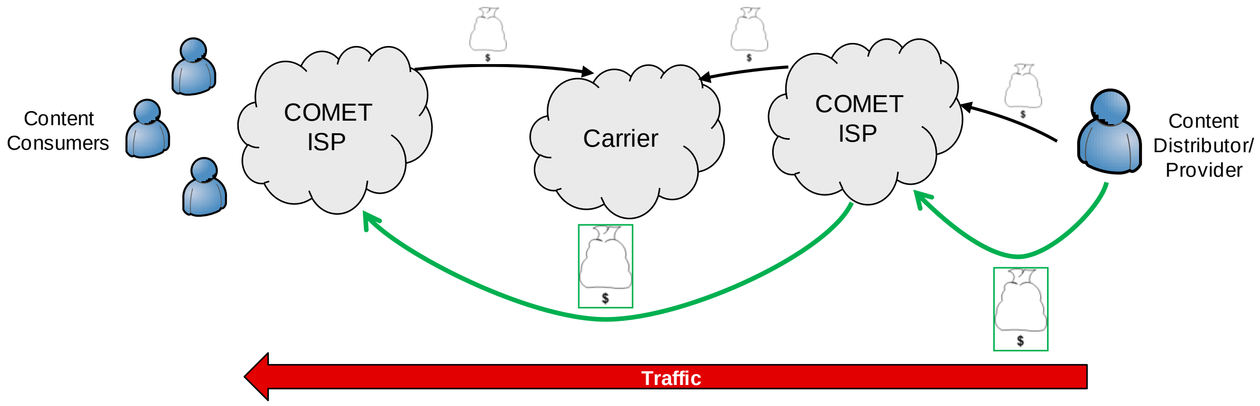

Figure 30: Scenario 1 for COMET business model related to point-to-multipoint communications

As it can be seen in Figure 30, there are two types of payments: those associated to data traffic and Internet connectivity (in black lines), and those related to agreements between edge COMET ISPs supporting point-to-multipoint content delivery through COMET (in green lines).

Regarding the payments related to data traffic, it is obvious that Carriers will receive lower revenues due to the decrease of traffic between edge COMET ISPs. Content Providers' ISPs also have lower incomes in terms of data traffic due to the reduction of the necessary bandwidth from their local Content Distributors/providers.

On the other hand, Content Providers' ISPs are expected to capture new incomes from Content Distributors/Providers thanks to price-competitive offers. Content Providers will be willing to reach global audiences, while reducing their bandwidth expenses. Content Providers' ISPs will share part of their revenues with Content Consumers' ISPs, with whom they will establish agreements related to point-to-multipoint content delivery.

Thanks to the COMET system, both edge domains can be aware of the content requests and perform an accounting of the active flows. In this way, charging between edge ISPs regarding content delivery can be done on a content basis according to the number of content replicas in the destination ISP, since now both ISPs can contrast their figures of content flows.

Although this scenario presents clear advantages over the current state of the art, it has scalability problems due to the numerous agreements to be established between ISPs, establishing also pSLSs with a mesh of ISPs. Next figure shows such situation. The number of agreements could be reduced if the number of Content Providers' ISPs is reduced to a small set, but this strongly depends on the dynamics of that market evolution.

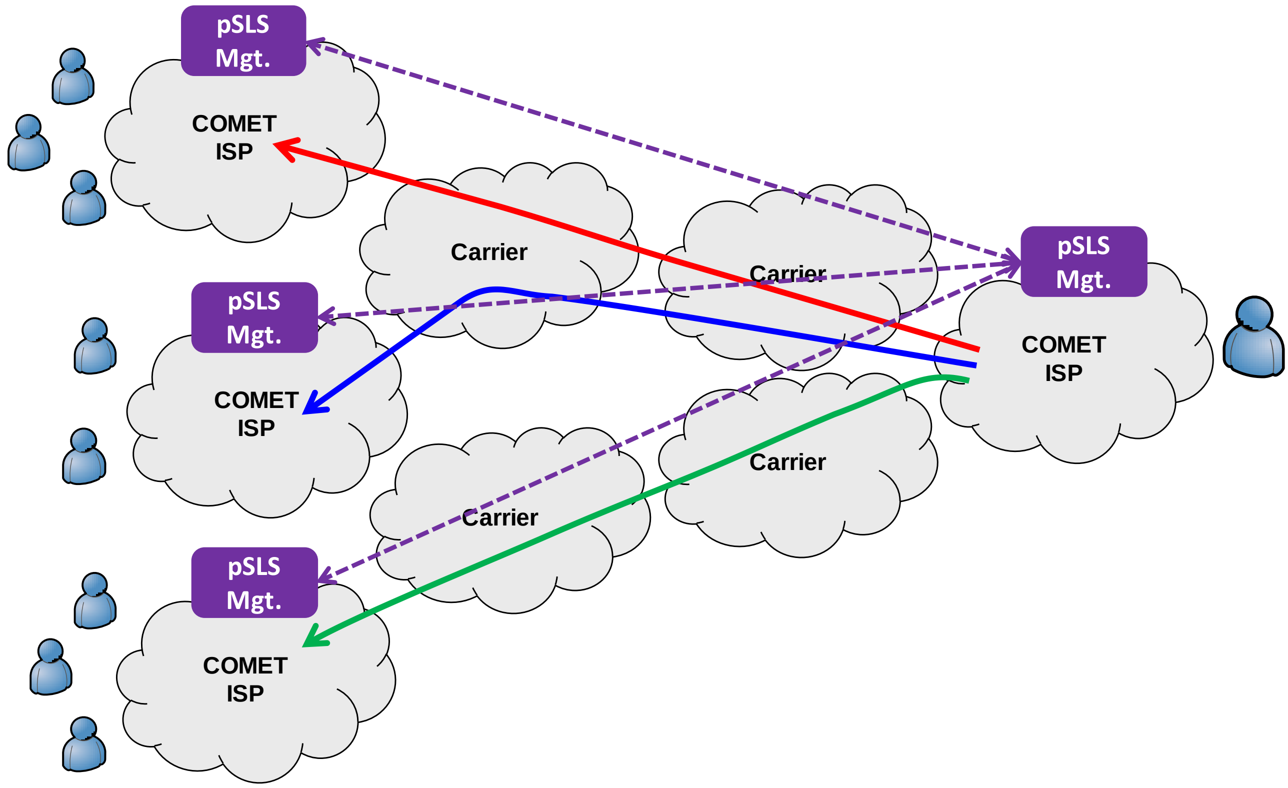

Figure 31: Scalability issues regarding interconnection agreements in the

<u>Scenario 2: COMET-aware Carriers supporting point-to-multipoint content delivery</u>

The natural evolution of the previous scenario is a scenario where Carriers support point-to-multipoint content delivery through COMET and are aware of COMET content requests, so that it is possible to perform charging in terms of the flows generated per content in each domain. The business model related to this scenario is shown in Figure 32, where cascading payments appear from Content Distributors/Providers to consumers' COMET ISPs through the whole chain of COMET Network Operators. These cascading payments would capture the costs incurred by each Network Operator according to the number of flows per content.

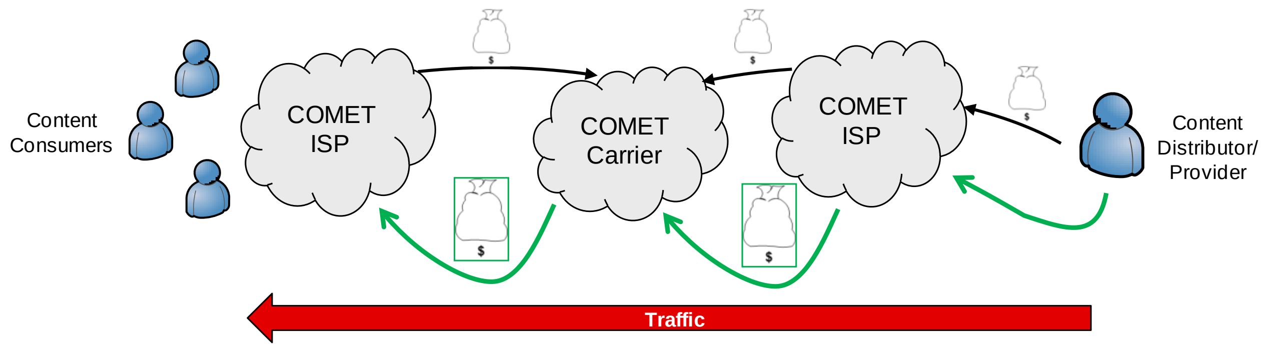

Figure 32: Scenario 2 for COMET business model related to point-to-multipoint communications

Carriers receive fewer revenues for delivering traffic, although now they are able to capture new incomes from ISPs, who receive new incomes from Content Distributors/Providers. Each COMET Network Operator will establish pSLS with adjacent Network Operator, in a similar way as in the inter-domain QoS cascaded model described in section 5.5.1 of this document.

5.6 Other business considerations in COMET

5.6.1 Impact of COMET in content publication

Regarding content publication, the following cases in the current Internet architecture can be identified:

- *Content Providers* own the content (e.g. CNN) and provide it directly to their customers (end-users).
- *Content Creators* produce the content and are the actual owners of the content. *Content Creators* might be either large organizations, whose content is their intellectual property or single end-users. In any case, a *Content Creator* has to search for a relevant *Content Provider* who will publish it to the Internet. Single end-users creating their own contens will have to either pay *Content Providers* for publishing their contents or offer them for free with the minimum required infrastructure, while large *Content Creators* charge with fees *Content Providers* for the delivery of their intellectual property to end users.

In both cases, any *Content Provider* has the authority and the infrastructure to publish contents in the Internet. Particularly, we can always assume that *Content Providers* own servers where contents are stored (content servers) taking into account parameters such as content popularity, content proximity to end-users, etc.

A *prosumer* aiming at publishing a content of his own must use a Content Provider interface (e.g. YouTube publication web page) or a *Content Distributor* interface (e.g. a tracker's interface) to register his content along with all the required metadata and get in return an identifier (e.g. URL, torrent descriptor, etc.) with which the content becomes addressable in the Internet (this identifier might also be publicized via e-mail, advertisements, etc). This procedure may be repeated as many times as the number of available distribution technologies. Next figure shows this publication procedure.

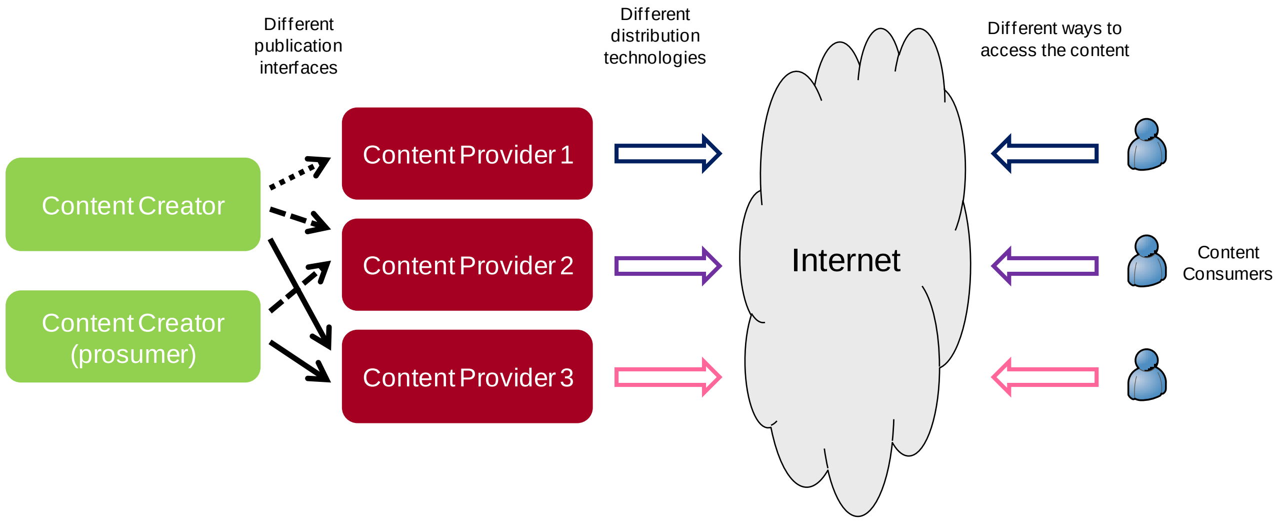

Figure 33: Content publication and access in today's Internet

On the other hand, the COMET system will introduce a unified platform for content publication (see Figure 34), in which *Content Providers* have the opportunity to store their contents in their respective content servers and publish their indexes directly to the COMET system, independently of its way of distribution. Besides, *prosumers* generating their own content would avoid contacting a hosting provider or creating a torrent (as they would do nowadays) in order to publish their content. Using the COMET system, such users could even store their contents in their local hard-disk drives, publish them directly to COMET and act like content servers when a user requests the specific content. Consequently, possible fees and subscriptions to hosting providers and torrent sites can be avoided, giving additional incentives to end-users and large organizations to share contents.

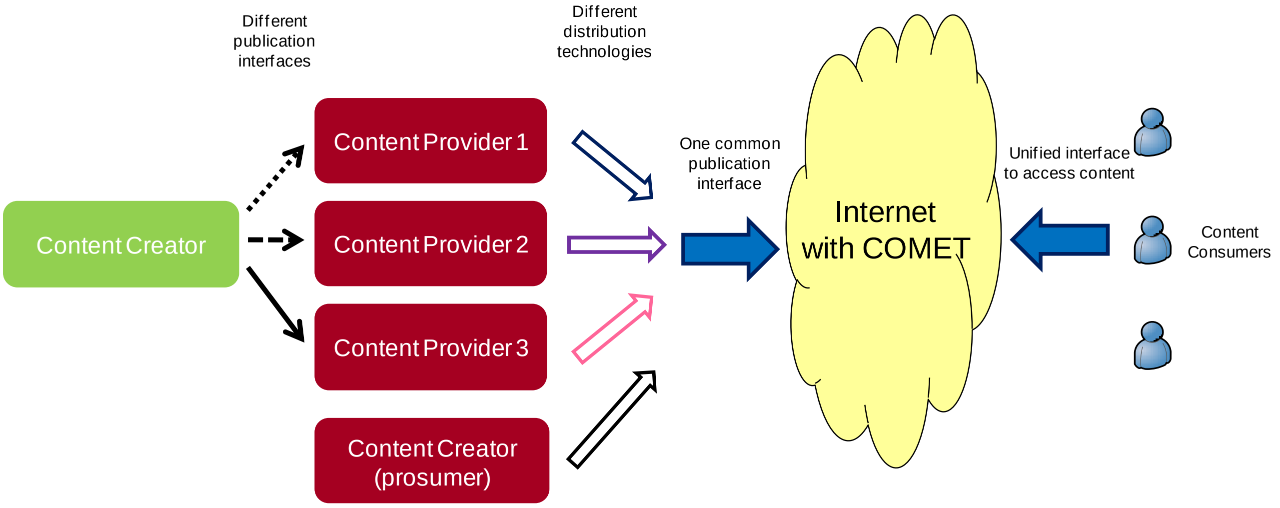

Figure 34: Content publication and access with COMET

It must be noted that the adoption of COMET could bring some changes regarding the payment relations related to publication. Part of the charges that were previously paid to the *Content Providers* may now have to be transferred to the *Network Operators* deploying the COMET system. Hence, a slight increase of connection charges might be expected as the *Content Creators* would be able to access a content-aware network. By offering this new type of content-aware "services", *Network Operators* adopting the COMET system would get advantage over their competitors and possibly increase their customer base.

5.6.2 Interface to third-party search engines

The capabilities of the COMET system as a mediator for content publication and consumption allow COMET ISPs to have more information about the content, for instance:

- An identifier of the content creator, which is also the owner of the content
- A description of the content, introduced during the publication process
- Properties of the content, such as the text and audio languages or the video codecs
- Keywords introduced during the publication process as well
- The COMET ISP where the content was registered during the publication
- The location of the content severs
- The content popularity, registered with every content request through COMET

All the previous information would be worth for third-party search engines, who could exploit it in different ways:

- Using the keywords to improve content searching, thus providing better search results.
- Using the identifiers of the content creators to rank search results depending on the reliability of the content creator, or to rank in a higher position those content creators who pay for being more popular.
- Providing better search results based on locality of the content (it could be inferred from the COMET ISP where the content was registered) or the language.
- Providing better search results based on the popularity of the content.

Thus, the definition of an interface between ISPs and third-party search engines where this information might be exchanged for money, as part of an additional agreement, would be very convenient. Hence, additional business models could appear where ISPs might be paid by third-party search engines for content information gathered by the COMET system. Then, these third-party search engines would exploit the information to offer better results to the end-users, while charging corporations for targeted advertisement, and content creators for improving their ranking in search results. Figure 35 shows the entities involved in this business model and the payment relations (blue arrows show payment flows, while also identify SLAs between the different entities).

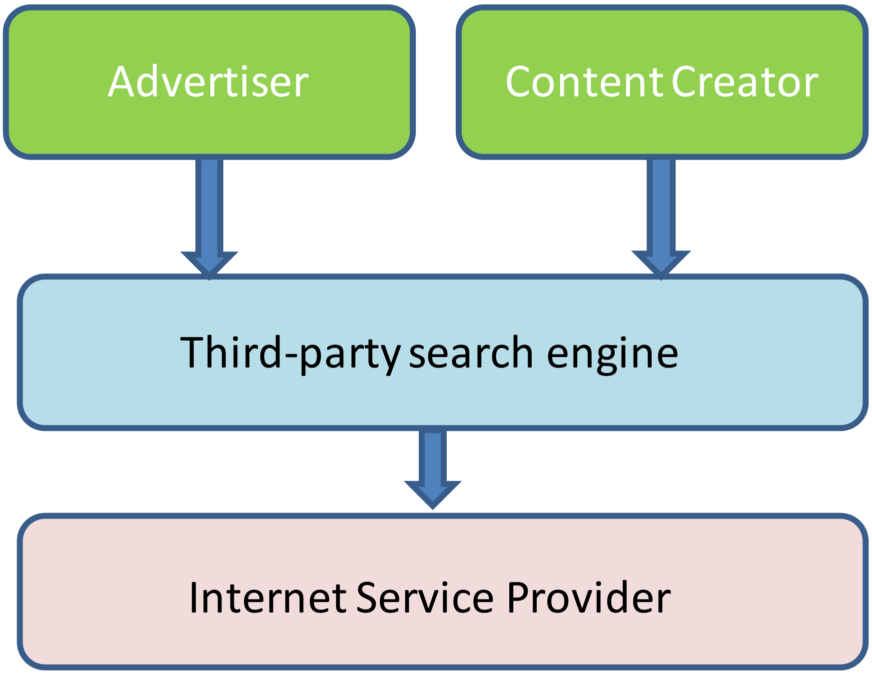

Figure 35: Business model related to the exchange of content-related information between ISPs and third-party search engines

6 Requirements for the COMET system

6.1 Introduction

This chapter presents the high-level system requirements that have been identified from the study of the different faces of the problem definition for content-aware networking. Those faces as sources of requirements are summarized in Figure 36:

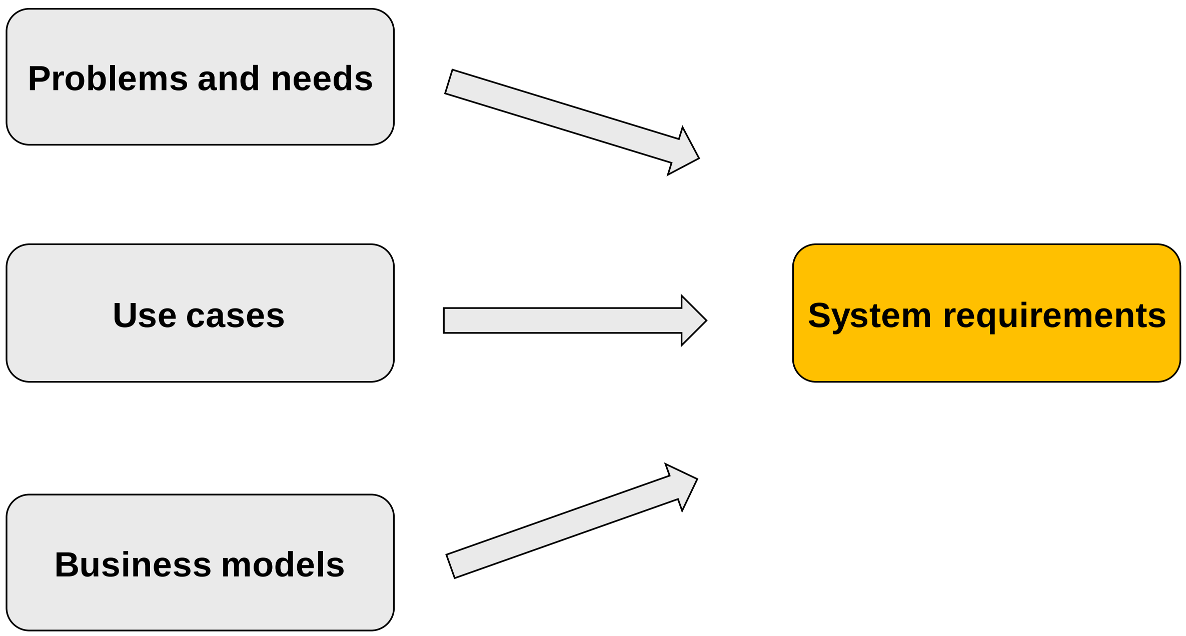

Figure 36: Conceptual diagram of the identification of system requirements

As it is shown in the diagram, as a result of the work done during the identification of the problems and needs existing in the area of content distribution, the search of use cases for COMET project and the study of the possible business models that could exist in content centric networks, the following requirements have been identified for the system.

6.2 System requirements

The identified system requirements are classified as follows:

a. Global requirements (Section 6.2.1),
b. Requirements for the Content Consumers (and Content clients) (Section 6.2.2.),
c. Requirements for the Content Providers (and Content servers) (Section 6.2.3.),
d. Requirements for the Content Mediation Plane (Section 6.2.4) and
e. Requirements for the Content Forwarding Plane (Section 6.2.5).

These system requirements will condition the design of the architecture of the COMET system.

6.2.1 Global requirements

- **The content must be treated as a primitive itself.** The architecture must be oriented to deal with all aspects of content natively, facilitating the access and distribution of contents. Support for safe, based on trusted content publication, friendly and fast content retrieval for consumers through the COMET architecture and mediation functionalities is required.

- **A global content naming and addressing scheme should be supported by an infrastructure capable of scalable content search and resolution.** The global content-aware mechanisms must be able to handle efficiently large amounts of content, being able to support significantly more objects than those handled by today's largest Content Distributors (YouTube, Flickr, Apple Store, for example). The protocols to be developed by the project will be capable of scaling to the order of billions (10^9) of content objects.
- **The COMET system should be open for future evolution of the Internet.** This can be achieved by the modularity in the design of different components and with a flexible high-level architecture.
- **Support for gradual and economical embracement of the COMET system by ISPs.** The designed architecture for content mediation and the associated mechanisms for content discovery, resolution and access must be scalable to be deployed in the largest ISPs, consisting of the order of hundreds of point of presence (PoPs) and core routers. These mechanisms and protocols should be applicable for content distribution at Internet-scale, involving autonomous networks of the order of tens of thousands of ASs.
- The content-aware mechanisms designed and developed for the network, when orchestrated by novel Content Mediation Plane (CMP) algorithms and protocols, **should facilitate the involvement of, potentially, all Internet users as Content Creators**. Thereby, creating the opportunity of a new, all-encompassing market where millions of small, medium and large Content Providers have access to efficient content distribution capabilities to reach billions of potential Content Consumers, taking advantage of a reduction of required resources, mainly bandwidth and processing capacity.
- The COMET system will support handover mechanisms which allow a **graceful switching of the content delivery path without impact on the application-layer**.

6.2.2 Requirements for the Content Consumers (and Content clients)

- **Access to the contents must be independent from the content location.** The naming architecture should guarantee location-independence, which in turn would guarantee smooth transition from today's host-centric to a future content-centric Internet.
- **The content identifier must be the same for different ways of distribution and nature of the content.** Also, different copies of content will be identified by the same Content-ID. It is, however, responsibility of the Content Providers to explicitly register the new copy of the content as such.
- **The Content Consumers must access the content in the same way as in current Internet** i.e. achieving user unawareness.
- The Content Client could optionally declare its capabilities during content resolution phase, but **it is up to the COMET system to decide how to deliver the content to the Content Consumer**.
- **The Content Client will obtain all the parameters necessary to invoke the application level requests**.

6.2.3 Requirements for the Content Providers (and Content servers)

- **There must be an interface that allows the Content Providers to update the content properties** (content location, server load, way of distribution, etc.)
- **The Content Provider should be able to establish policies** to enforce the way to publish and deliver the contents to the Content Consumers.

6.2.4 Requirements for the CMP (mediation layer requirements)

- **There must exist a global content resolution architecture** for efficient and scalable name and content resolution.
- **There should be an integrated traffic and resource management solution compatible with the content resolution architecture** to increase network efficiency and content delivery in order to reduce network congestion on the most highly loaded links.
- **There should be an information gathering system in the CMP** for collection of various performance metrics on networks and servers. This is going to be implemented in the COMET Monitoring Module.
- **The protocol interfaces between the CMP and the Content Providers, publishers and end user devices must be efficient.** The user terminals should be able to send their content consumption requests through these interfaces, and the Content Providers must announce their server condition and the information about the contents they publish using these interfaces. To complete this requirement, some others have been extracted from the use cases:
    - **The CMP must be able to dynamically modify the information related to the location of the servers in the content record.**
    - **The COMET system must offer to the Content Provider the possibility of registering different ways of distribution.**
- **The CMP in an ISP must be aware of network conditions** in order to take decisions oriented to reduce the latency in content retrieval that is due to network failures, network congestion or server load.
- **There should be some kind of interaction between the Content Mediation Plane and the Content Forwarding plane to enforce content delivery**.
- **The CMP, upon the content request from a user device, should be able to request capabilities to enhance or facilitate the QoS and multicast in the network** for the delivery of that content to that user device.

6.2.5 Requirements for the CFP (network layer requirements)

- **There must be a content forwarding architecture able to perform content-based forwarding at speeds similar to the ones in IP-based forwarding**.
- The elements in the CFP should support **QoS-aware content delivery**.
- The elements in the CFP should support **point-to-multipoint content delivery**.
- **Content may be cached in the network** to optimise network resource usage.
- **There should be an interaction between the CFP and the CMP to provide information on network conditions and, optionally, routing information**.

6.2.6 Summary table

The following table gathers the system requirements that have been explained in the previous chapters.

ID	Category	System requirement
1	**Global**	Content as a primitive
2	**Global**	Global content naming and addressing
3	**Global**	Open for future evolution of the Internet
4	**Global**	Scalable to be deployed in the largest ISPs
5	**Global**	Involvement of all Internet users as Content Creators
6	**Global**	Graceful switching of the content delivery path without impact on the application-layer
7	**Content Consumer**	Access independent from content location
8	**Content Consumer**	Content ID independent from way distribution and nature of content
9	**Content Consumer**	User unawareness
10	**Content Consumer**	Content Client able to declare his capabilities
11	**Content Consumer**	Content Client will obtain all necessary parameters
12	**Content Provider**	Interface to update the content properties
13	**Content Provider**	Capability of establishing policies to enforce the way to deliver contents
14	**CMP**	Global content resolution architecture
15	**CMP**	Integrated traffic and resource management solution to increase network efficiency and content delivery
16	**CMP**	Information gathering system
17	**CMP**	Efficient protocol interfaces
18	**CMP**	Capability of dynamically modify servers location information
19	**CMP**	Possibility of registering different ways of distribution
20	**CMP**	Network conditions and routing information awareness
21	**CMP**	Interaction between the Content Mediation Servers and the Content Aware Forwarders to enforce content delivery
22	**CMP**	CMP able to request the enforcement of QoS and multicast in the network
23	**CFP**	Content forwarding architecture able to reach IP-based forwarding speeds
24	**CFP**	Elements in CFP able to support QoS-aware content delivery
25	**CFP**	Elements in CFP able to support point-to-multipoint content delivery
26	**CFP**	Content may be cached to optimize network resource usage
27	**CFP**	Interaction between the CFP and the CMP to provide information on network conditions and, optionally, routing information

Table 1: System Requirements for the COMET system

7 Conclusions

The COMET system has been presented as a global architecture for the Internet which would allow *Network Operators*, and specifically *Internet Service Providers*, to act as mediators for content publication and distribution. This mediation would make possible an improvement of content delivery in terms of quality and effective bandwidth utilisation, which is extremely relevant for real-time events. Based on this mediation, the *ISPs* will be able to offer a unified interface for content access, making consumption simpler for the end-user and independent on the nature of the content, its location or its way of distribution. Besides, the COMET system will offer a publication interface for *Content Providers* so that every content registered in the COMET system will have a unique and global content identifier, bringing independence from the content location and way of distribution (e.g. HTTP transfer of P2P distribution).

Four use cases have been identified as drivers of the design of the architecture of the COMET system, and they will be used as starting point for future demonstration activities. The use cases are focused on the distribution of contents as well as on communication services where there could exist mediation by *Network Operators*, thanks to their content-awareness. The **main technical advantages** identified in the use cases are the following:

- Unified access to the content whatever its nature and location
- Content delivery with guaranteed QoS
- Point-to-multipoint content delivery capabilities, reducing bandwidth needs for live contents
- Graceful handover of the content delivery path, providing more resilience and flexibility for multi-homed users
- Advanced publication mechanisms, allowing *Content Providers* to update content servers on-the-fly, while switching among different ways of distribution.

Based on the previous benefits, the deliverable illustrates the possibilities for the **COMET system** to encompass former and new business models related to content distribution, **becoming a flexible umbrella for all the variety of content-related business that are either present in today's Internet or those which would be potentially possible in an evolved Internet, with advanced content-handling features**.

The business models based on free content access provided currently by Over-The-Top distribution systems can be easily adopted into the COMET framework. With the adoption of the COMET system, *Network Operators* would become active part of a content-oriented Internet. Enhanced content access will be offered by *ISPs* to *Content Consumers*, while *Content Providers* and *Distributors* would be benefited from the network mediation offered by *ISPs*.

Besides, new business models based on charged content access (either in a pay-per-view scenario or in a subscription-based scenario) are possible with the COMET system. On the one hand, with COMET redefining content delivery, current Internet-based broadcasting services, such as VoD or Live streaming services, would benefit from an improved delivery both in terms of flexibility and performance. For instance, *Content Providers* would be able to charge *Content Consumers* for the delivery of high definition real-time content with QoS guarantees, being able to reach wider audiences for live content with affordable bandwidth costs. Moreover, *prosumers* generating their own content would not require to contact a hosting provider (as they do nowadays) in order to publish their contents. Using the COMET system, *prosumers* could even store their contents in their local hard-disk drives and publish them directly to the COMET system acting as *Content Providers* on user requests. In this case, *ISPs* would become a more active part of the value chain for content delivery, being even able to charge *Content Providers* on a content basis (and not just in terms of traffic volume).

On the other hand, thanks to the ease of content publication through the COMET system, *ISPs* would have the ability to open their IPTV platforms to third-party *Content Providers*, following a revenue sharing model, similar to the one being followed by some ISPs to export other capabilities (e.g. OpenTelefónica, BT Web21C API, etc.). The relation between the *Content Providers* and the

IPTV platform would involve an agreement to show the published content and the payment from the IPTV platforms to the *Content Providers* according to the actual number of consumers. This scenario will bring benefits to all agents:

- *ISPs* would offer a larger and more attractive catalogue of channels through their IPTV platforms, being even able to incorporate long-tail contents.
- *Content Consumers* would benefit from a larger catalogue and would have more flexibility to configure their channel catalogue.
- *Content Providers* from Internet will be able to reach larger and more consolidated audiences through the ISP's IPTV platforms. Moreover, small or minority *Content Creators* could offer their contents through the IPTV platform, requiring much less infrastructure and expenses.

Additional business models could appear for search engines, which might be willing to pay for content information gathered by the COMET system. These third-party search engines would exploit the information to offer better results to the end-users, while potentially charging corporations for targeted advertisement, and content creators for improving their ranking in search results.

Finally, as a result of that detailed analysis on use cases and business models, the deliverable identifies **27 high-level requirements for the COMET system**, encompassing global requirements as well as specific requirement for the content clients, content servers, the CMP and the CFP. These system requirements capture requirements derived from:

- the problems and needs existing in the area of content distribution that are intended to be solved with the COMET system,
- the technical advantages highlighted in the use cases, and
- the identified business models.

These high-level system requirements are currently driving the design of the COMET architecture.

8 References

[1] Krzysztof Walkowiak, QoS Dynamic Routing in Content Delivery Networks, NETWORKING 2005: 1120-1132

[2] Jaison Paul Mulerikkal, An Architecture for Distributed Content Delivery Network. A minor thesis submitted in partial fulfilment of the requirements for the degree of Masters of Applied Science (Information Technology).

[3] Content Delivery Networking by IPTV Operators Paves the Way for High Quality Internet Video, Article by Alcatel-Lucent http://www.alcatel-lucent.com/enrich/v2i32008/article_c4a3.html

[4] http://www.pressreleasepoint.com/frost-amp-sullivan-video -cdn-market-grow-15-billion-2013-telcos-gain-stakes

[5] Al-Mukaddim Khan Pathan, R. B. (2007). *A Taxonomy and Survey of Content Delivery Networks*. Australia: Grid Computing and Distributed Systems Laboratory, The University of Melbourne.

[6] George Pallis, A. V. (2006). *Insight and Perspectives for Content Delivery Networks*. ACM.

[7] Kangasharju, J., Roberts, J., & Keith, R. (2002). *Object replication strategies in content distribution networks*.

[8] Sjöberg, D. (2008). *Content Delivery Networks: Ensuring quality of experience in streaming media applications*. TeliaSonera International Carrier

[9] A. Asgari, M. Boucadair, D. Griffin, P. Georgatsos and G. Pavlou, "Interconnection Models for QoS-Based IP Service Offering", White Paper, to be published

[10] M. Howarth, P. Flegkas, G. Pavlou, N. Wang, P. Trimintzios, D. Griffin, J. Griem, M. Boucadair, P. Morand, H. Asgari, P. Georgatsos, "Provisioning for Inter-domain Quality of Service: the MESCAL Approach", IEEE Communications, Vol. 43, No. 6, pp. 129-137, June 2005

[11] ATT ICDS, http://www.att.com/icds

[12] CDNlist.com, http://www.cdnlist.com

[13] CDN Market Growth 2006 – 2009, http://www.researchandmarkets.com/reportinfo.asp?report_id=615197

[14] P. Mockapetris, "Domain Names – Concepts and Facilities". Internet Engineering Task Force (IETF) Request for Comments (RFC), RFC 1034, November 1987. Available: http://www.ietf.org/rfc/rfc1034.txt

[15] P. Mockapetris, "Domain Names – Implementation and Specification". Internet Engineering Task Force (IETF) Request for Comments (RFC), RFC 1035, November 1987. Available: http://www.ietf.org/rfc/rfc1035.txt

[16] Cisco Visual Networking Index: Forecast and Methodology, 2009-2014, Cisco White Paper, Cisco Inc., 2010. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf

[17] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, October 2009.

[18] EU FP7 SmoothIT project, URL: http://www.smoothit.org/

[19] The MBone Deployment: http://tools.ietf.org/wg/mboned/

[20] The Internet2 project: http://www.internet2.edu/

[21] What Is A Webinar?, ASME, http://www.asme.org/Education/Courses/Webinars/Webinar.cfm

[22] Red5 Open Source Flash Server, http://osflash.org/red5

[23] Rajkumar Buyya, Mukaddim Pathan, "Content Delivery Networks: Overlay Networks for Scaling and Enhancing the Web", Chennai, India 2008

[24] N. Wang, D. Griffin, J. Spencer, J. Griem, J. Rodríguez Sánchez, M. Boucadair, E. Mykoniati, B. Quoitin, M. Howarth, G. Pavlou, A. J. Elizondo, M. L. García Osma, P. Georgatsos and O. Bonaventure, "A framework for lightweight QoS provisioning: network planes and parallel Internets" Proc. IFIP/IEEE Int. Symposium on Integrated Network Management (IM 2007), short paper, Munich, Germany, 21-25 May 2007.

[25] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weis, "An Architecture for Differentiated Services", IETF RFC 2475, December 1998

[26] IST project TEQUILA, http://www.ist-tequila.org

[27] IST project MESCAL, http://www.mescal.org/

[28] IST project EUQOS, http://www.euqos.eu/

[29] GSMA IP eXchange model, http://www.gsmworld.com/our-work/programmes-and-initiatives/ip-networking/ip_exchange.htm

[30] Akamai Technology, http://www.akamai.com/html/technology/index.html

[31] J. He et al, "Towards Internet-wide Multipath routing", IEEE Network, Vol. 22, Issue 2, pp.16-21, 2008

[32] Z. Li et al, "QRoN: QoS-aware Routing in Overlay Networks", IEEE JSAC Vol. 22, pp29-40, 2003

[33] OpenMeetings, http://code.google.com/p/openmeetings/

[34] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, L. Wei, IETF RFC2362 "Protocol Independent Multicast-Sparse Mode (PIM-SM)"

9 Abbreviations

AGAVE	A liGhtweight Approach for Viable End-to-end IP-based QoS Services
AS	Autonomous System
BGP	Border Gateway Protocol
BW	BandWidth
CAF	Content Aware Forwarder
CDN	Content Delivery Network
CFP	Content Forwarding Plane
COMET	COntent Mediator architecture for content-aware nETworks
CMP	Content Mediation Plane
CMS	Content Mediation System
CPU	Central Processing Unit
DHT	Distributed Hash Table
DNS	Domain Name System
DoW	Description of Work
DSL	Digital Subscriber Line
FLV	FLash Video
HD	High Definition
IBP	Internet Backbone Provider
ICDS	Intelligent Content Distribution Service
IP	Internet Protocol
IPTV	Internet Protocol TeleVision
IPX	Internetwork Packet eXchange
ISP	Internet Service Provider
LSP	Label Switch Path
MH	Multi-homing
MIME	Multipurpose Internet Mail Extensions
OTT	Over-The-Top
P2P	Peer-to-Peer
PC	Personal Computer
PCE	Path Computation Element
PoP	Point of Presence
pSLS	provider-level Service Level Specification
QC	QoS capability
QoC	Quality of Content
QoE	Quality of Experience
QoS	Quality of Service

RPC	Remote Procedure Call
SLA	Service Level Agreement
SmoothIT	Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies
STREP	Specific Targeted REsearch Project
TCP	Transmission Control Protocol
TV	TeleVision
URL	Uniform Resource Locator
VoD	Video on Demand
VoIP	Voice over Internet Protocol

10 Acknowledgements

This deliverable was made possible due to the large and open help of the WP2 team of the COMET project within this STREP, which includes besides the deliverable authors as indicated in the document control. Many thanks to all of them.

Annex A: COMET terminology

Carriers	Those Network Operators which are interconnected with each other and form the backbone of the Internet. They have large capacity networks and their main responsibility is the delivery of the content to other Carriers or ISPs (wholesale services).
Content Delivery Network (CDN)	A networked system of computers interconnected through which cooperate —in a transparent way to the users— in the efficient and predictable distribution of contents based on geographic location, the origin content and the content delivery server.
Content	The actual piece of data that the Content Consumer requests (web-page, file, video stream etc.).
Content aware Forwarder (CAF)	The entity responsible for enabling content distribution across the global Internet.
Content Client	The actual end host machine which is the destination of a content flow.
Content Consumer	The entity that consumes (e.g., watches, listens to, or reads) the content.
Content Consumption	The process initiated by a Content Consumer to receive the requested content.
Content Creator	The entity that that owns the rights of the content and wants to publish it to the Internet. It is the author of the content.
Content Distributor	The entity that owns and maintains infrastructure that supports the distribution of content.
Content Forwarding Plane (CFP)	One of the two COMET planes responsible for locating the actual content server that hosts the requested content or the best content server candidate if there are multiple available.
Content Identifier (ID)	An opaque (i.e. non human-readable) string that is used by COMET system in order to point to the specific content requested by the Content Consumer.
Content Mediation Plane (CMP)	One of the two COMET planes in charge of the delivery of the content once requests are resolved based on its current knowledge of both the network and server status.
Content Mediation Server (CMS)	The entity responsible for content manipulating, for instance content publication, resolution and delivery operations.
Content Name	A human-readable string that explicitly and uniquely identifies the requested Content.
Content Provider	The entity that specializes in storing and making content available to the Content Consumers.
Content Publication	The process of making content available to Content Consumers.
Content Resolution	The process responsible for discovering the requested content based on the given content ID.
Content Server	The server that actually hosts content is referred to as Content Server.
Internet Backbone Providers (IBPs)	See "Carriers".

Internet Service Providers (ISPs)	Those Network Operators which have some network of their own, usually country-wide. They may serve a number of individual customers (home users or companies). Their main responsibility is the delivery of the content to the Content Consumers (retail services).
Metadata	A set of parameters related to a specific content which are given to COMET system and identify the content properties required for accessing a piece of content (e.g. content servers, ways of distribution, QoS requirements…).
Name resolution	The process responsible for translating the human-readable content name to a machine-oriented content ID.
Network Operator	The entity that provides networking services, wired or wireless.
Over-the-top (OTT)	Content distribution system in which Content Providers have capabilities to go directly to Content Consumers with their contents, bypassing traditional network gatekeepers and access providers, that is, not allowing the Network Operators they go through to take an active role or obtain any revenue from the content distribution.
Prosumer	An end-user that is at the same time a Content Consumer and a Content Provider.
Web-inar	A presentation, lecture, workshop or seminar that is transmitted over the Web. A key feature of a Web-inar is its interactive elements: the ability to give, receive and discuss information.