COntent Mediator architecture for content-aware nETworks

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

Deliverable D6.1
Demonstration Scenarios and Test Plan

The COMET Consortium

Telefónica Investigación y Desarrollo, TID, Spain
University College London, UCL, United Kingdom
University of Surrey, UniS, United Kingdom
PT PLC, PT, 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. David Flórez Rodríguez
Telefónica Investigación y Desarrollo, dflorez@tid.es
Don Ramón de la Cruz
28043 Madrid
Spain

Phone: +34 91 312 88 42
E-mail: comet@comet-project.org

Document Control

Title: Demonstration Scenarios and Test Plan

Type: Public

Editor(s): David Flórez, Francisco Javier Hernández Romero

E-mail: dflorez@tid.es, fjhr@tid.es

Author(s): Wei Koong Chai (UCL), Ioannis Psaras (UCL), George Pavlou (UCL), David Flórez (TID), Francisco Javier Hernández Romero (TID), Andrzej Beben (WUT), Jordi Mongay Batalla (WUT), Jarek Sliwinski (WUT), George Petropoulos (ICOM), Spiros Spirou (ICOM), Michael Georgiades (PT)

Doc ID: 20120621-D61-v3.5-FINAL.docx

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
V0.1	October 10th	David Flórez	Release of draft ToC
V0.2	November 4th	David Flórez	First Contributions to Sections 3 and 4
V0.3	November 21th	David Florez, Javier Hernández	Contributions to functional test (5.1) and performance SNME test (5.2.3)
V0.4	November 24th	Andrzej Beben, Jordi Mongay Batalla	Contribution about local multicast use case, tests related to RAE and CAFE, comments about federated testbed description
V0.5	December 2nd	George Petropoulos	CME and CRE test cases
V0.6	December 5th	Andrzej Beben, Jordi Mongay Batalla	Functional Tests fot multicast to client
V1.0	December 7th	David Flórez	First Draft
V1.1	December 21st	Wei Koong Chai	Description of Comet-Emvision interface
V2.0	December 22nd	David Flórez	Final Version Ready for release
V3.0	April 24th	David Flórez	First Version of revised performance tests to address reviewrs' comments
V3.2	June 14th	David Flórez, George Petropoulos, Sergios Souros, Andrzej Beben,	Second version of revised performace tests to address reviewers' comments
V3.3	June 15th	David Flórez	Prefinal Version
V3.4	June 18th	Wei Koong Chai Michael Georgiades	Internal Review
V3.5	June 22th	David Flórez	FInal Version ready for resubmission

Table of Contents

(This page is left blank intentionally.)

1 Executive Summary

In this deliverable we define the guidelines that will rule the demonstration and testing activities for COMET System in WP6, covering the following tasks:

- Specification of the COMET's federated testbed to be set up between TID's, WUT's and PT's premises in order to emulate a real situation where several ISPs (Internet Service Providers) are involved.
    - This testbed will be used only for demonstrating the Decoupled Approach and not the Coupled one, given that the Decoupled approach is an evolutionary approach that can be deployed over the current Internet and therefore the existing facilities in those organisations can be easily adapted for Decoupled Approach's purposes.
- Detailed description of the Use Cases to be demonstrated over the federated testbed and the procedures/operations required to support them.
- Detailed description of the individual tests to be carried out, both functional and performance-oriented, whose results will be gathered in deliverable D6.2

Section 2 of this document provides a brief introduction of WP6 efforts, while section 3 provides an overview of the layout of the Federated COMET Testbed (i.e., which COMET entities are implemented and how many of them will be deployed on each testbed according to the proposed scenarios and available resources at each site). Therefore, although each partner offering an individual testbed site will emulate an ISP, each of the involved partner has distinct aspects of the COMET approach that are highlighted in the testbed, WUT will then create three different ones in their premises, in order to emulate a transit ISP during the retrieval of contents and to allow for more complicated paths between CCs (Content Client) and CSs (Content Servers). PT will host the hierarchy of CREs (Content Resolution Entity) to illustrate the content resolution process. Finally, TID will be more centred in deploying and hosting CCs, CSs and SNMEs (Server and Network Monitorin Entity).

Section 4 deals with the description of the Use Cases that are going to be demonstrated over the federated testbed. From the four cases defined in D2.1 [2] and D2.2 [2], two have been selected to be demonstrated, because they are deemed to be the most suitable for validating the requirements and functionalities of COMET. The Use Cases are:

- *Case 1, Adaptable and efficient content distribution*, that will cover most of the capabilities and functionalities of the COMET System, Route Awareness, Server Awareness, Content Publication and Content Comsumption
- *Case 4, P2P Offloading*, demonstrating how the COMET System can offload the load to P2P (Peer to Peer) CSs in the case where all the main distribution ones are overloaded, avoiding denial of service in case of peaks affecting content requests for consumption.

The operations carried out by COMET in each use case are described over the testbed layout defined in Section 3; this shows how COMET would behave in an almost real scenario. Besides, for the use case 1, a subcase has been defined to demonstrate how COMET can incorporate and manage multicast services in the Client's ISP, instead of the Server's ISP, avoiding resource exhaustion if too many remote end-users request the multicast content.

Section 5 enumerates the test cases defined for validation of the system. They can be grouped in two main categories:

- Functional tests, which check in detail the different operations carried out by COMET for each use case.
- Performance tests, which characterise each of the COMET architectural elements. CME (Content Mediation Entity), CRE, RAE (Routing Awareness Entity), SNME, and CAFE (Content-Aware Forwarding Entity) in order to measure, among others, medium and peak response times and amount of queries the modules can manage simultaneously.

Scalability issues are not dealt with in this deliverable, they will be reported in WP5's deliverables.

Section 6 enumerates, wherever possible and pertinent, how the requirements defined in WP2 have been satisfied by the tests defined in this deliverable. Special attention is given to the performance metrics that are gathered in D5.1 [6].

Finally, Section 7 focuses on the description of the envisaged COMET-ENVISION integration. For this integration, ENVISION (Co-optimisation of overlay applications and underlying networks) will supply COMET with cost path information over the links leading from one ISP to their neighbours, helping the CME to refine its CS/Path decision process. This means that ENVISION's CINA (Collaboration Interface between Network and Application) server will have to be deployed in the tesbed's ISP and an interface will have to be implemented to connect CINA to the CME.

2 Introduction

The main target of WP6 activities is the testing and demonstration of the COMET System, as designed, implemented and integrated in previous work packages (WP2,WP3, WP4, WP5), over a testing ground as close as possible to a real Internet situation. Therefore, this deliverable will define the testing and demonstration activities whose results will be compiled in D6.2, mainly for the COMET Decoupled approach only since this can be directly deployed over the current Internet and , and hense therefore no disruptive changes/modifications will be required in the already existing testbeds made available by COMET's partners.

To enhance the similarity of the envisaged test ground to a real world situation, the layout of the testbed relies on several basic concepts:

- The setup can emulate several interconnected ISPs, as in the real Internet.
- An end user can request contents located in a different ISP from home ISP, so that typically an end user's ISP, a CS's ISP and one or more intermediate ISPs will be involved in the resolution/retrieval of a content.
- Users can retrieve contents hosted in ISPs other than theirs.
- The same content can be distributed by different ISPs.
- There can be multiple network paths allowing the retrieval of a specific content from a CS by a CC. These paths will be differentiated by assigning different QoS/BW constraints to them.

Having these requirements in mind, the aim of this deliverable is threefold.

First, The description of the federated testbed connecting the individual testbeds located at partners' premises (TID, WUT and PT), which will emulate at least three different ISPs interconnected by multiple paths. This federated testbed is described in Section 3, restricted to an overview of the general rules and features that will govern the testbed setup and deployment, leaving the final details for D6.2 that will summarise the results of demonstration activities

Secondly, the selection of the Use Cases to be demonstrated over the federated testbed from the set defined in D2.1 [2] and D2.2 [3], as well as a complete explanation of how these uses cases can be demonstrated on the envisioned federated testbed, which will be the focus in Section 4.

Thirdly, the definition of a complete set of test cases, gathered in Section 5, with focus on two different aspects:

- Functional tests, which will prove that the Use Cases operate as intended,
- Performance tests that will characterise the behaviour of the different architectural entities (CME, CRE, SNME, RAE, CAFE) in terms of response time and/or maximum amount of queries that can be managed.

Scalability issues are not dealt with in this deliverable. They will be reported in WP5's deliverables.

In Section 6, the qualitative requirements defined in D2.2 [3] and the performance metrics defined in D5.1 [6] are mapped to the test cases addressing them (if pertinent and possible), so that it is shown that most of COMET capabilities and functionalities are covered by the testing and demonstration activities in WP6.

Finally, Section 7 describes how COMET's and ENVISION's functional elements are going to be integrated and the repercussions it will have on the federated testbed.

3 Description of COMET Federated Testbed

The aim of this section is the description of the layout of the COMET Federated Testbed where the Decoupled Approach will be tested and demonstrated.

The federated testbed will consist of three sites, located at three partners' premises (namely WUT, TID and PT), which will be interconnected by means of IPv6 on IPv4 tunnels, as depicted in the following figure.

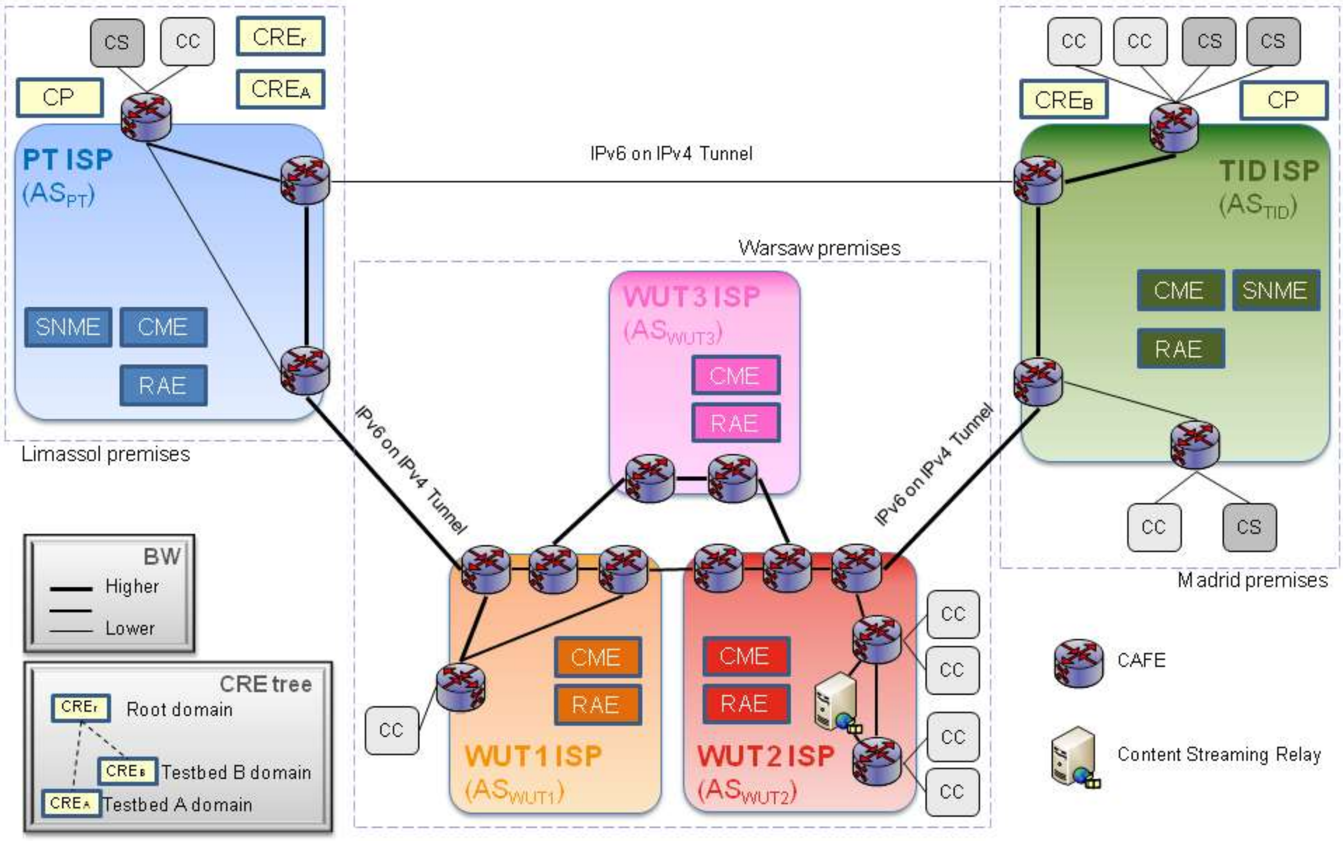

Figure 1: Envisaged Layout of the Federated Testbed

It is important to note that the description of the local testbed and their elements is not final. This is because more elements (especially CS, CC and even CAFEs or CREs) can be deployed according to the needs arising from integration tasks in WP5 and testing activities in WP6. The final description of the federated testbed will be provided in D6.2. Nevertheless, the layout illustrated above includes the minimum set of necessary entities each testbed should contain, according to the resources and expertise of each partner.

Thus, given below, these are the rules and conditions that will govern the final federated testbed. Unless explicitly stated otherwise, the OS of the machines hosting the elements deployed in the testbed is Linux.

- Each local testbed will stand for a local ISP. Therefore, PT's testbed will be an ISP identified as AS_{pt} (AS stands for Autonomous System) and TID's testbed an ISP identified as AS_{TID}. However, in order to allow for a more complex ISP scenario closer to reality (i.e. where data flows could cross several interconnected ISPs), WUT's testbed will host three different ISPs, identified as AS_{WUT1}, AS_{WUT2} and AS_{WUT3}, interconnected between them.

- In the decoupled approach, the CREs are not associated with a specific ISP (hence the term decoupled). There is a hierarchy defined among the CREs, similar to DNS, from a root CRE to those CREs serving contents located in one or several domains. In the COMET testbed, this is emulated by placing the CRE root at PT's testbed, as well as a local a CRE, and another local CRE at TID's premises. No CRE will be deployed at WUT's premises. Hence, any content distributed from a future CS located within WUT's testbed will have to be managed by an external CRE (i.e. either PT's or TID's CRE).
- Each ISP will mandatorily host a CME and a RAE. The RAE will dicover the paths from this ISP to the remaining ones. The CME will resolve a Content Name into the address and the features of a server distributing the content via communication with the CRE hierarchy, as well as configuring the best path from the distributing CS to the requesting CC.
- A number of CCs will be deployed in each ISP in order to enable the retrieval of the contents mediated by COMET. The final amount of CCs will depend on the use casesUse Cases described in section 4 and the tests sketched in section 5.
    - For testing purposes, the clients will be machines running Windows OS.
- CSs serving the contents mediated by COMET will also be deployed in the testbeds. The number of CSs and the type of contents being distributed will vary according to the use casesUse Cases described in Section 4 and the tests sketched in Section 5.
    - As a rule of thumb, most of CSs will be located at TID's and PT's premises, which will act as source testbeds WUT will be mainly allocated with the task to demonstrate how COMET can cope with multiple ISPs.
    - The CSs will serve contents in three different basic ways: Streaming, VoD and P2P.
        - Streaming and VoD will be implemented by using VLC [7].
        - P2P services will be implemented by using microTorrent [8].
    - For testing purposes the servers will be machines running Windows.
- In those the ISPs where CSs have been deployed, a SNME will also be deployed to gather the status of those CSs and attend queries regarding their status from the CME.
- Any ISP containing CSs will also include a Content Publisher in order to allow the publication of Content Records in the CREs.
- A number of CAFEs will be deployed inside each testbed. The CAFEs act as edge or border routers.
    - The edge CAFEs are located at edges of access networks, where the CSs and CCs are connected. Several CSs and CCs can be served by the same CAFE and the population of CSs and CCs can be segmented in sets to be managed by different CAFEs, as it will likely happen in a real situation (and as it have been devised for TID's testbed as depicted in Figure 1).
    - The border CAFEs are located at the endpoints of ISPs, where inter-domain links are connected. We use IPv6 over IPv4 tunneling between border CAFEs located in different premises, and IPv6 over Ethernet to connect ISP domains located in the same geographical location (i.e., WUT premises).
    - The connectivity between CAFEs located inside a given ISP is provided by IPv6 over (VLAN) Ethernet.
- Each ISP will support COMET CoSs on either intra- or inter-ISP connections that will differ in provisioned values of QoS/bandwidth parameters. The COMET system uses this information to choose the best path/server according to the content transfer requirements and the CC's CoS. In Figure 1, these different path characteristics are being represented by different traces (the better the QoS/BW, the wider the trace).

   - Though not dynamically, the provisioning of particular domains could be changed in order to change the configuration of routing paths inside the testbed and test the Route Awareness, Path Discovery and Decision Process capabilities of COMET.
   - The CCs inside an ISP will be assigned different CoS, according to the IP address they use to connect to the testbed. These differences in the CoS will be used for testing the path discovery and decision capabilities of COMET.
- One of WUT's domains will include the Content Streaming Relay (CSR) entity. The CSR supports local multicast in the client's domain. It receives content streams from border CAFE and sends unicast or multicast streams to a number of CCs.

4 Description of Use Cases

This section explains how the four Use Cases defined in D2.1 [2] and D2.2 [3] can be tested in the COMET federated testbed.

For demonstration purposes of the decoupled approach, the following two Use Cases have been selected out of the initially considered four cases:

- **Use Case 1: Adaptable and efficient content distribution.** This use case will demonstrate the basic functionality of the COMET System in full detail. In other words:
    - How a Content Owner can publish contents in COMET.
    - How COMET is able to resolve the Content Name assigned to a published content and return the characteristics of a CS distributing that content.
    - How COMET can create an end to end path from the selected CS to CC with assured QoS/BW according to the client's CoS and selected CS.
    - How a Content Owner can add more CSs to distribute a content if needed, and how these new servers will be automatically and dynamically taken into account by the decision process
    - How a Content Owner can distribute the same content by different means (i.e application and transport protocols). In other words, the Content Owner can deploy different types of CS, for Streaming, VoD, P2P or direct download, under the umbrella of a single Content Name/Content Identifier.
    - How end users need not be aware of changes of the CSs' distribution protocols (or its population or location). In other words, the Content Name originally assigned to the content will be valid, regardless on the changes of the CSs characteristics, and COMET will be able to assign a server for content retrieval.
    - How local multicast can be provided by Content Streaming Relay.
- **Use Case 4: P2P Offloading.** This use case will demonstrate how COMET can switch from assigning higher quality sources (i.e. Streaming) to lower quality sources (i.e., P2P), in the case where the high quality sources are about to be overloaded, preventing a disruption in the QoS of users already retrieving the content from the high quality server.

The rationale behind short-listing the Use Cases enumerated in D2.1 and D2.2 from four to two is as follows

- Use Case 1 will test most of the features and requirements sketched in D2.1 [2] and D2.2 [3]. In that sense this is the most complete use case and the one where demonstration effort should be focused.
- **Use Case 2 (Handover of content delivery path in a multi-homing scenario)** is more related to the management and interworking of different access technologies (fixed, wireless, mobile) than to pure COMET concepts. Therefore, and in order to avoid a federated testbed excessively complicated, this use case will not be demonstrated in the federated testbed, in spite of its intrinsic interest and utility.
- **Use Case 3 (Webinar "All about CDNs")** involves the creation and management of meta Content Records consisting of several independent Content Records. Since most of the control and synchronisation of the different Contents compiled by the Meta Content Record will be carried out by an external webinar server, it does not seem as the most suitable use case for demonstrating the decision and routing capabilities of COMET.
- User Case 4 complements Use Case 1 by building on it but in this case demonstrating how the COMET system can cope with contents simultaneously distributed by different sources, with different protocols and QoS/BW requirements (in this case VoD vs. P2P)

Futhermore, for Use Case 1 a subcase has been defined in order to test and demonstrate a multicast solution in the Client's ISP, as opposed to multicast from the server.

The following sections explain in greater detail how these Uses Cases can be demonstrated on the envisaged federated testbed.

4.1 Adaptable and efficient content distribution (Use Case 1)

From the point of view of Content Owners and end users, this use case consists of the following stages:

- A Content Owner wants to distribute a live event (i.e. a football match). The Content Owner publishes a Content Record containing a number of Content Sources describing the characteristics and requirements of the streaming servers distributing the Content.
- An end user requests a content identified by the Content Name assigned to the previous Content Record. COMET will return the connection parameters (protocols, Server's IP, etc) of a server distributing the content, thus enabling the downloading of the Content.
- The Content Owner can add as many new streaming servers as needed (i.e. if the audience exceeds the initial projection) to the Content Record created for the content. end users requesting this Content Name will be assigned to the fresh streaming servers, according to the server's load and QoS/BW constraints.
- Once the live event has ended, the Content Owner will remove the streaming servers Content Sources describing this and add new Content Sources for distributing the content by means of VoD servers. These changes will be applied on the Content Record originally created, neither creating a new one nor changing the Content Name.
- New users requesting the same Content Name will be served by VoD server after this.

4.1.1 Prerequisites

Before accepting users' requests, a number of COMET entities have to be configured.

First, each CME has to gather the information about the paths leading from its local ISP in to the rest of ISPs. For instance, AS_{WUT1} will learn the paths to AS_{WUT2}, AS_{WUT3}, AS_{TID}, AS_{pt}. This operation (called Routing Awareness) will be performed by the RAEs deployed in each ISP, which will feed the CME with the routes discovered so far. These paths consist of the AS numbers of the ISPs from ISP source to ISP destination and are qualified with QoS parameters (packet loss and delay) and BW, so that these characteristics allow each CME to find out the most suitable path to a CS when a Content Name is to be retrieved.

The second phase to be completed before testing can start is the Content Publication process. The Content Owner will use the Content Publisher to create a Content Record in the CRE for each content to be distributed and assign a Content Name to it. Typically, a Content Record consists of a number of Content Sources that group CSs with common characteristics (full description can be found in Deliverable D3.2[4]). The features of interest in the selection of the final path between the end user and the CS are briefly sketched below.

- Requirements of quality (BW and QoS) that affect the final path chosen by the CME, because the paths gathered by the RAE are characterized by the same parameters.
- The COMET Class of Service hints at the SLA subscribed by the end user, and can take the values Premium (PR), Better than Best Effort (BTBE) and Best Effort (BE), so that the Content Owner can group CSs in different categories and allocate those with better performance/resources for the end users who have subscribed for higher CoS. Therefore, implicitly PR>BTBE> BE.
- Priority Flags allow establishing a hierarchy among Content Sources with the same CoS, so that a path solution for the Content Sources with the higher priority will be first tried and, if no solution is found, lower priorities will be used in turn. For instance, a Content Owner could deploy Content Sources for VoD and Direct Download, and decide that the end user should be directed first to VoD sources and then to Direct Download ones, when the former is overloaded. In such case, the priority should be used to implement this decision rule.

4.1.2 Description

End users (CCs) and CSs can be located anywhere in the Federated Testbed, and in fact, for testing purposes, they will be deployed and tested in multiple different locations. However, for the sake of the explanation, it will be considered hereafter that the CCs are located in one of the ISP under WUT umbrella (AS_{WUT1}, AS_{WUT2} and AS_{WUT3}), while the CSs (and thus the Contents) are deployed either at TID's (AS_{TID}) or PT's testbeds (AS_{PT}).

The first step requires that RAEs gather the information about the paths linking the different ISPs. Once those paths have been correctly propagated to the CMEs, a streaming server will be activated (i.e. launching VLC in Streaming mode) at PT's and at TID's testbeds. At the same time a Content Record containing a single Content Source with information related to both CSs and their connection parameters will be published in the CRE by using the Content Publisher. For this example, the CoS will be defined as Premium (more complex examples with different Sources, QoS (Quality of Service)/BW(Bandwidth) requirements and CoS are given in section 4.1.3).

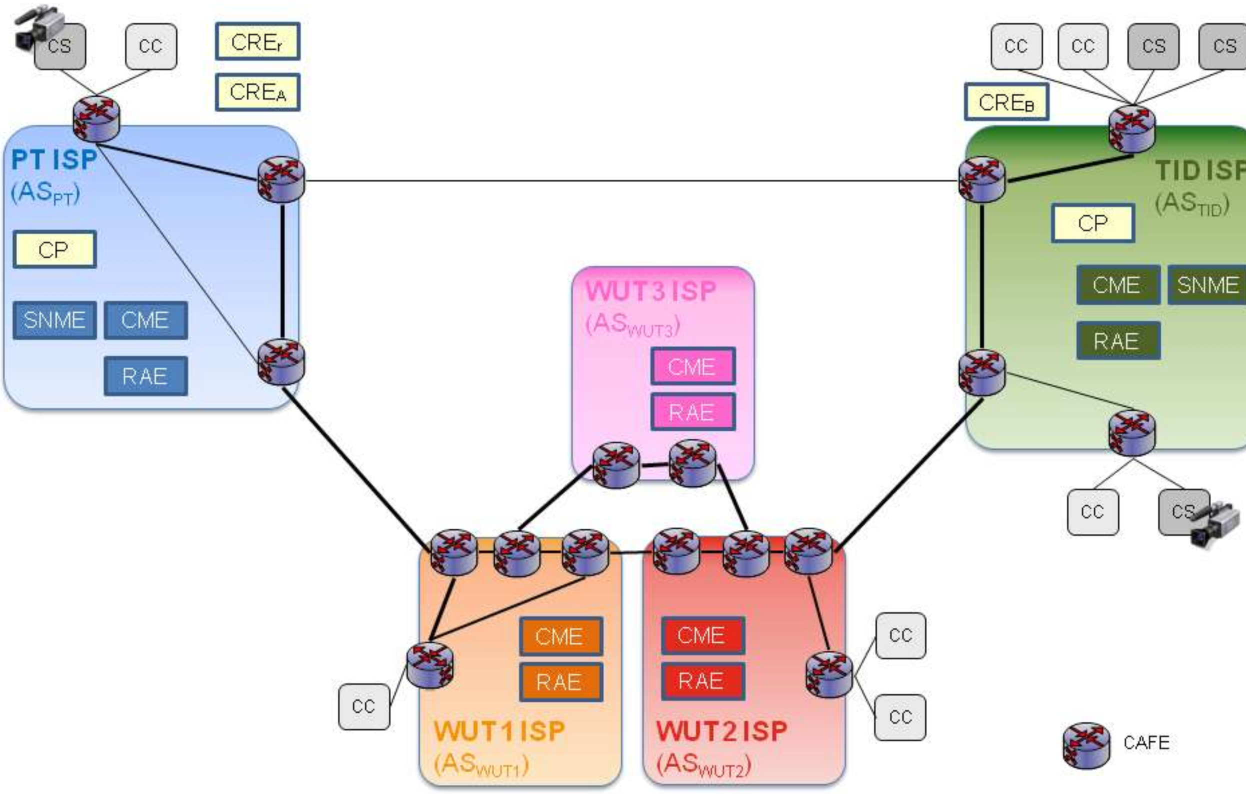

Figure 2: Initial Situation after CSs' Activation and Publication

The following steps are taken:

1. An end user (located in AS_{WUT1}) gets to know of the Content Name needed to retrieve the content and requests it using the Content Name via Web Browser. Automatically, the CC is invoked (because the Browser identifies the Content Name as a COMET URL) and the CME is requested to provide the characteristics of the CS delivering the content identified by this Content Name.
2. The CME retrieves the Content Record associated to this Content Name from the CRE. In this case, the Content Record will contain the information about a single Content Source, which describes two different CSs, one at TID's and another at PT's premises. It is assumed that the Client's CoS and the Content Source's CoS are the same and set to Premium.
3. The CME checks the local information about the paths that lead to the ISPs hosting the CSs enumerated in the Content Source.

4. The client's CME sends a query to each CME of the server's ISPs in order to retrieve the status of the CSs as provided by remote SNMEs (in this case the SNMEs will report that there are no overloaded servers).
5. The client's CME assesses the collected information and decides to choose the path between the AS_{WUT1} and AS_{PT}. This is because the route to the CS at AS_{TID} is of worse quality than that leading to AS_{WUT1}, as illustrated in Figure 2.
6. The selected path needs to be configured and the client's CME (AS_{WUT1}) contacts to server's CME (AS_{PT}) to translate the path expressed as a list of AS numbers into a path expressed as a list of keys from the CAFE attached to the CS to the CAFE attached to the CC.
7. Server's CMEs configure the server's attached CAFE with the list of CAFEs required to get from one to the other.
8. The client's CME sends to the CC the required information for connecting to the selected CS.
9. Finally, the client downloads the content through the selected path.

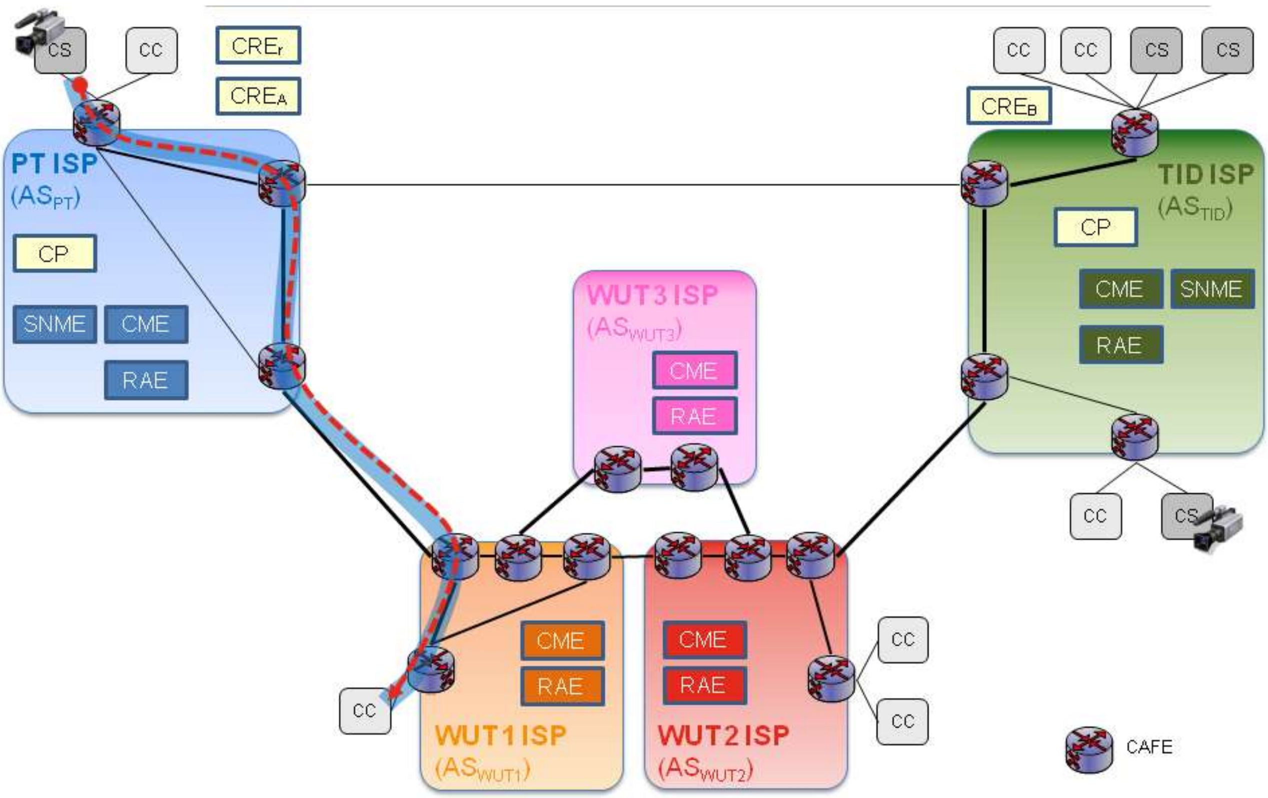

Figure 3: Path Setup between CC at AS_{WUT1} and a CS at AS_{PT}

In case the load of the servers increases, requiring more servers to satisfy the increment in viewers, new streaming CSs can be added to the existing Content Record by the Content Owner. In the example, a new streaming server will be activated in AS_{TID}. The list of servers in the associated Content Source will be updated with the connection parameters of the new server.

When a new end user (this time located at AS_{WUT2}) requests the same Content Name, the sequence of steps explained above is repeated. In this case, the path to the new CS at AS_{TID} is probably of better quality than that leading to AS_{PT} , because of the number of hops, and the request will be directed to the CS at AS_{TID} (the old CS at AS_{TID} is already in a path of worse quality, so it will not be taken into account in the decision process). If both paths were of the same quality, the switch between servers could be simulated by increasing the load of the CS at AS_{PT}, forcing that the decision depends on the information provided by the SNMEs.

The following picture shows the new path configured:

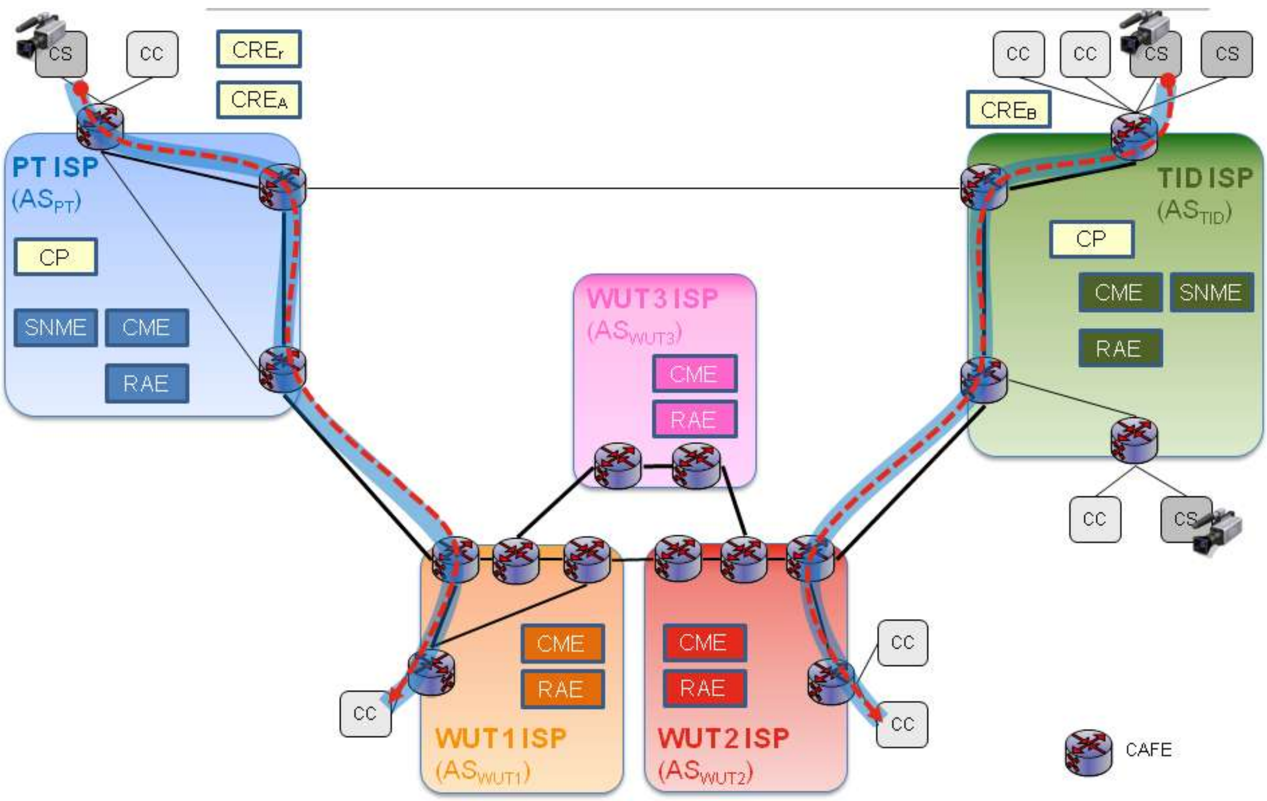

Figure 4: Path Setup between Client at AS_{WUT2} and a CS at AS_{TID}

Once the live event has finished, the streaming servers stop broadcasting, terminating the streams to the end users. If the Content Owner wants to make available a recording of the event, it will have to perform the following tasks:

- To activate a CS that will provide a VoD service (in this case VLC will be launched in VoD mode). In our example, two servers of this sort are deployed at AS_{TID}.
- The Content Source for the servers providing the Streaming Service has to be deleted from the Content Record at the CRE.
- A new Content Source with the requirements and characteristics of the VoD servers has to be added to the existing Content Record.

When an end user (in this case located at AS_{PT}), requests the content with this Content Name, COMET follows the same steps described above, but the path will be established with one of the two new CSs.

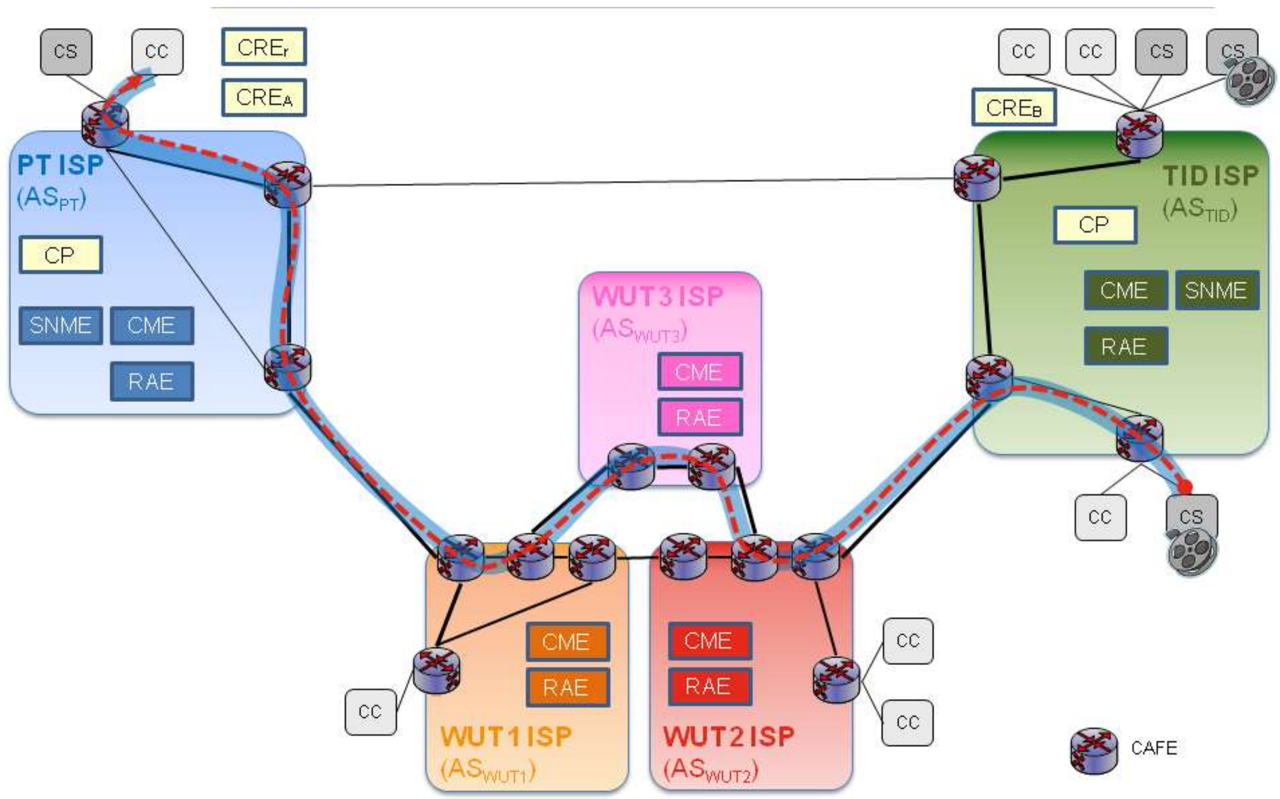

Figure 5: Example of path in the Case of VoD CSs

4.1.3 More Complex Scenarios

The paths that RAEs fed into the CME are qualified in terms of QoS/BW. This information is considered as long term and will be normally defined when the individual links between CAFEs are provisioned. As shown in Figure 1, the COMET federated testbed can be preconfigured to assign specific weights to each link between a pair of CAFEs, thus creating a complex network of paths in terms of QoS/BW.

Taking this into account, different Streaming/VoD CSs can be deployed over the entire testbed, and such Content Record can hold a variety of Content Sources that will implement:

- Different CoS
- Different BW/QoS requirements.

If a number of CCs with different CoSs (the CoS for a CC is extracted from its IP, though in a real system a more complex query to a provisioning system would be required) are deployed in different locations across the federated testbe, it can be checked that:

- Only the Content Sources allowed for a client's CoS are used in the decision process. In general, a user with CoS BE cannot use BTBE or PR Content Sources, while a user with CoS BTBE cannot use PR Content Sources.
- The Path is assigned according to the BW/QoS Requirements defined in the Content Source, implying that if a path to the CSs defined in the Content Source that matches the Client's CoS is not found, an error will be returned and other Content Sources should be tried in turn, if available.
- For CSs in a Content Source whose paths are the same or indistinguishable in terms of QoS, the decision will be governed by the information provided by the SNME according to the load of each individual CS.

4.1.4 Point to multipoint streaming

This scenario shows distribution of live streaming content supported by Content Streaming Relay (CSR). The CSR is designed to support point to multipoint disctibution of live streaming content related to known known to be popular a priori. The CSR receives a single content stream from CS and provides it to a number of CCs in the client's domain. The exemplary scenario is presented in Figure 6. This scenario assumes that the CS is located in TID domain and CSR is located in WUT2 domain.

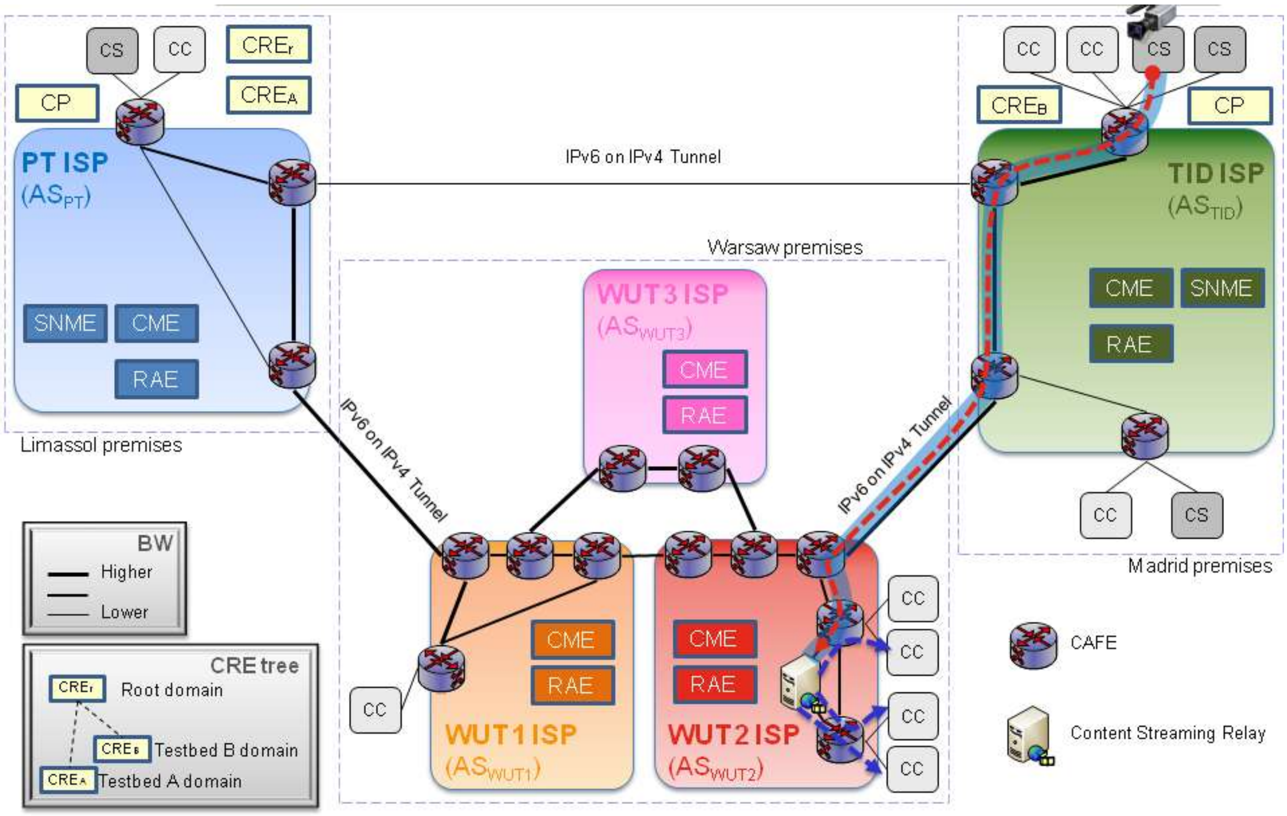

Figure 6: Example of local multicast in Client's domain

The point to multipoint streaming requires the following steps:

1. In the case of planned popular events, the operator of WUT2 ISP prepares the CSR. The operator uses CC to initiate content resolution procedure and prepare content delivery path going from the best CS to the CSR. Next, the operator runs the CSR using the obtained URL.
2. The CSR becomes an additional local source of the content, which would be used to serve users located in the client's domain. Therefore, the operator runs the SNME to monitor the status of CSR and registers the CSR in the COMET system as new CS. In this scenario, we assumed that CSR is registered in Content Record, so the operator in consultation with content owner updates Content Record by adding new CS. At this stage, the CSR is ready to serve users requests.
3. An end user located in AS_{WUT2} writes the Content Name in its Web Browser. Then, the CC is automatically invoked and it sends request to the CME.
4. The CME retrieves the Content Record associated to the Content Name in order to get information about the content and the available content sources. As long as the CSR is not overloaded, it will be selected to serve users in the client's domain.
5. The CME sends to the CC the information required for connecting to the CSR.
6. Finally the following users' requests will be directed to the CSR.

In this scenario, we demonstrate the capabilities of COMET system to reduce bandwidth consumption compared to the content delivery models used in the current Internet.

4.2 P2P Offloading (Use Case 4)

From the point of view of Content Owners and end users this use case consists of the following stages:

- A Content Owner wants to distribute a content and publish it. This results in the creation of a Content Record corresponding to this new content describing the characteristics and requirements of the streaming servers distributing the content.
- An end user requests the published content, which is identified by the Content Name assigned to the previous Content Record. COMET will return the connection parameters (protocols, server's IP address) of a server distributing the content, thus enabling the downloading of the Content. To this regard, this case would be similar to Use Case 1.
- However, the Content Owner also deploys CSs with P2P capabilities and adds a Content Source in the same Content Record already created with the description of these P2P servers, in order to offload the traffic to the P2P Servers if the streaming ones become overloaded.
- This P2P Content Source will be assigned a lower priority than the Streaming Content Source's, forcing the COMET system to give priority to an available streaming server first. If the decision process does not return a solution, P2P Sources are tried next.

4.2.1 Prerequisites

As with Use Case 1, the RAEs in the different ISPs have to propagate the path information and store them in their adjacent CMEs.

From the point of view of content publication, a single Content Record for the content is to be created in the CRE, consisting of two Content Sources, one for the streaming servers, and another for the P2P Servers. We note that:

- CoS should be the same for both Content Sources, because the service is targeted to the same type of users. P2P servers are treated as fallback for the case where the streaming services will not be able to provide a service of the required quality if more users are allowed to download the content. However, the QoS/Bandwith requirements for the Streaming Source should be of better quality than those for the P2P Source to ensure a higher level of QoS for the intended primary means of delivery. In other words, streaming servers are considered as higher quality sources while P2P as lower quality sources.
- The Streaming Source has to be assigned higher priority than the P2P Source, so that the P2P servers are only used when the streaming servers have become overloaded as notified by the SNME monitoring them.

4.2.2 Description

End users (CCs) and CSs can be located anywhere in the COMET Federated Testbed, and in fact, for testing purposes, they will be deployed and tested in multiple different locations. However, for the sake of clarity, it will be considered hereafter that the CCs will be located in one of the ISPs under WUT umbrella (AS_{WUT1}, AS_{WUT2} or AS_{WUT3}), while the CSs (and thus the contents) of both types will be deployed at TID's (AS_{TID}).

The first step consists of the RAE gathering the information about the paths linking the different ISPs. Once those paths have been correctly propagated to the CMEs, two streaming servers will be activated (launching VLC in Streaming mode) at TID's testbed, as well as a P2P one (launching a microTorrent in tracker/seed mode). A single Content Record containing two Content Sources with information related to both types of CSs and their connection parameters will be published in the CRE by using the Content Publisher.

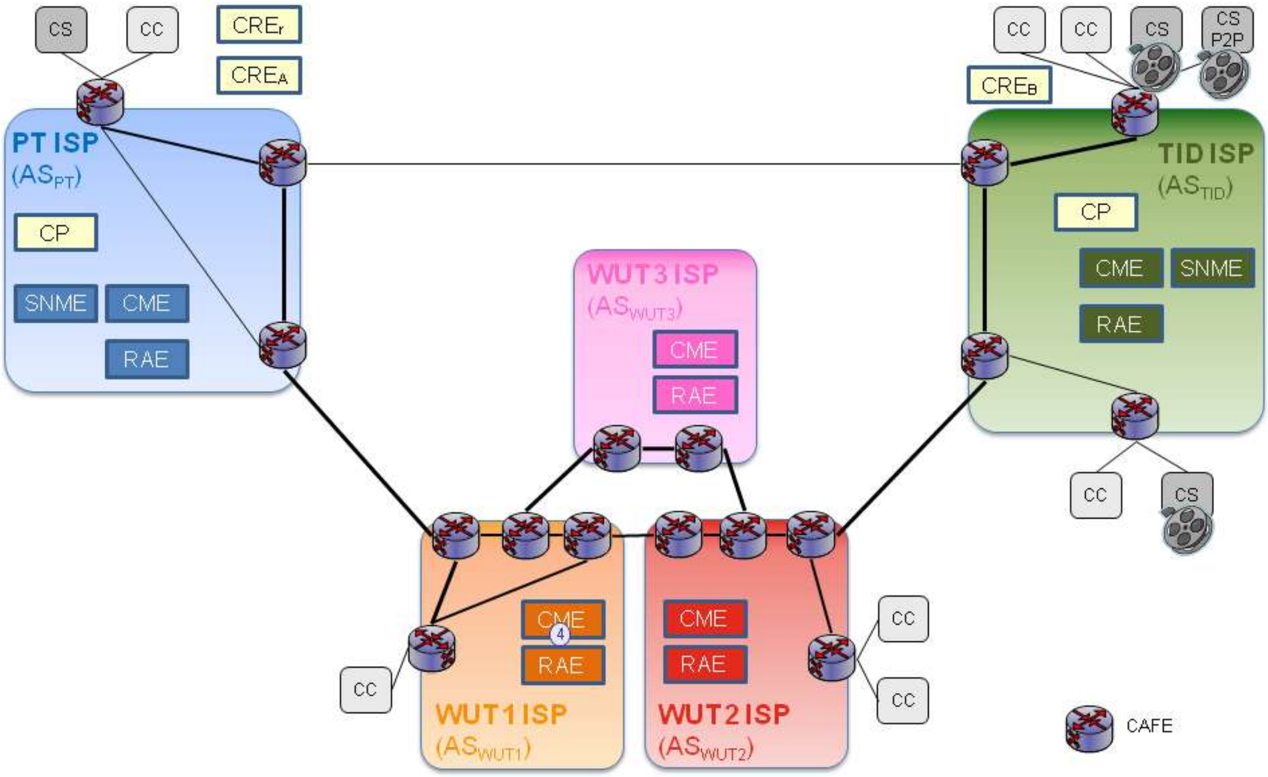

Figure 7: Initial Situation for P2P Offload

Initially, when an end user requests the content via the Content Name (by writing in the Web Browser) the steps followed will be the same as in Use Case 1. The content will be downloaded from one of the streaming servers, once the tunnel has been established, as depicted in the following figure.

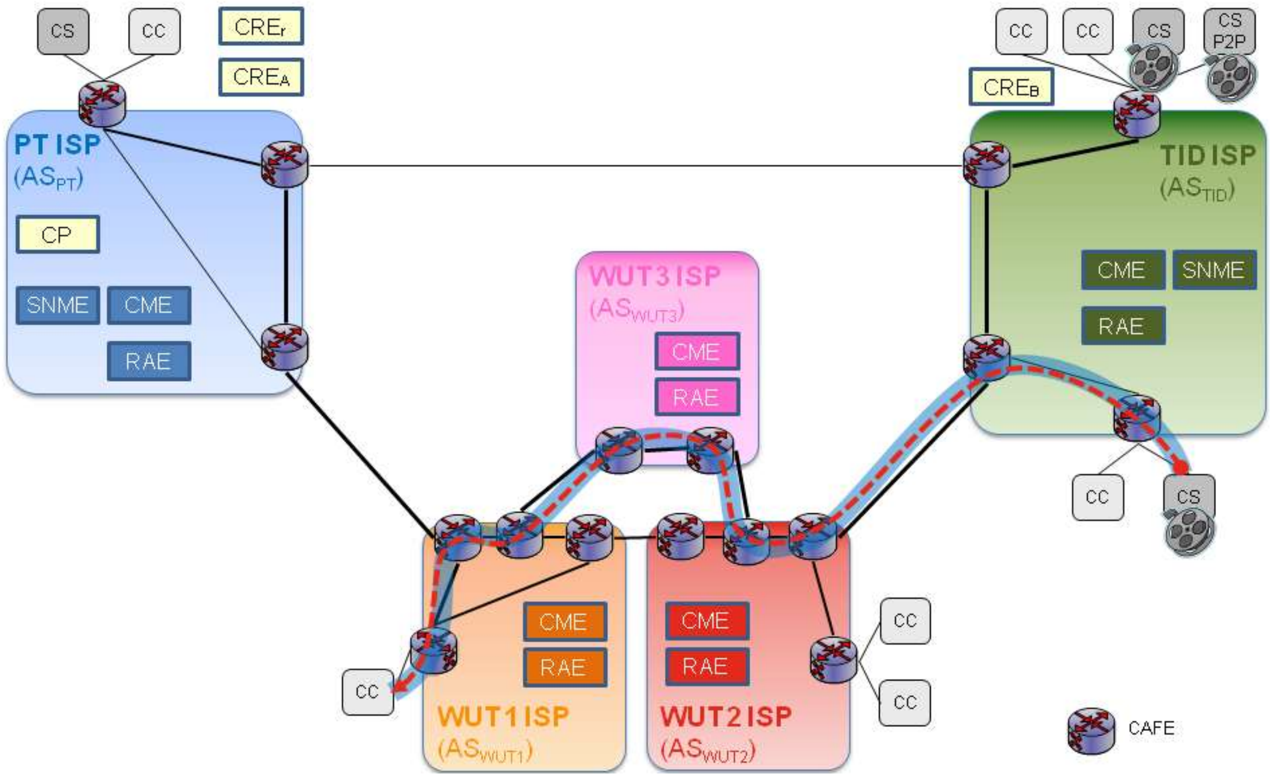

Figure 8: Path Established with a Streaming CS

The main difference with Use Case 1 is when the streaming servers are too overloaded to accept new clients. This will be notified by the SNMEs monitoring their statuses. When this happens, if a CC, located for instance at AS_{WUT2}, requests the retrieval of the content identified by the Content Name, the CME will not find a solution for the Streaming Sources and will fall back on the P2P Sources, leading to the assignment of a P2P CS for the downloading of the content and the establishment of a path between both ends, as depicted in the following figure.

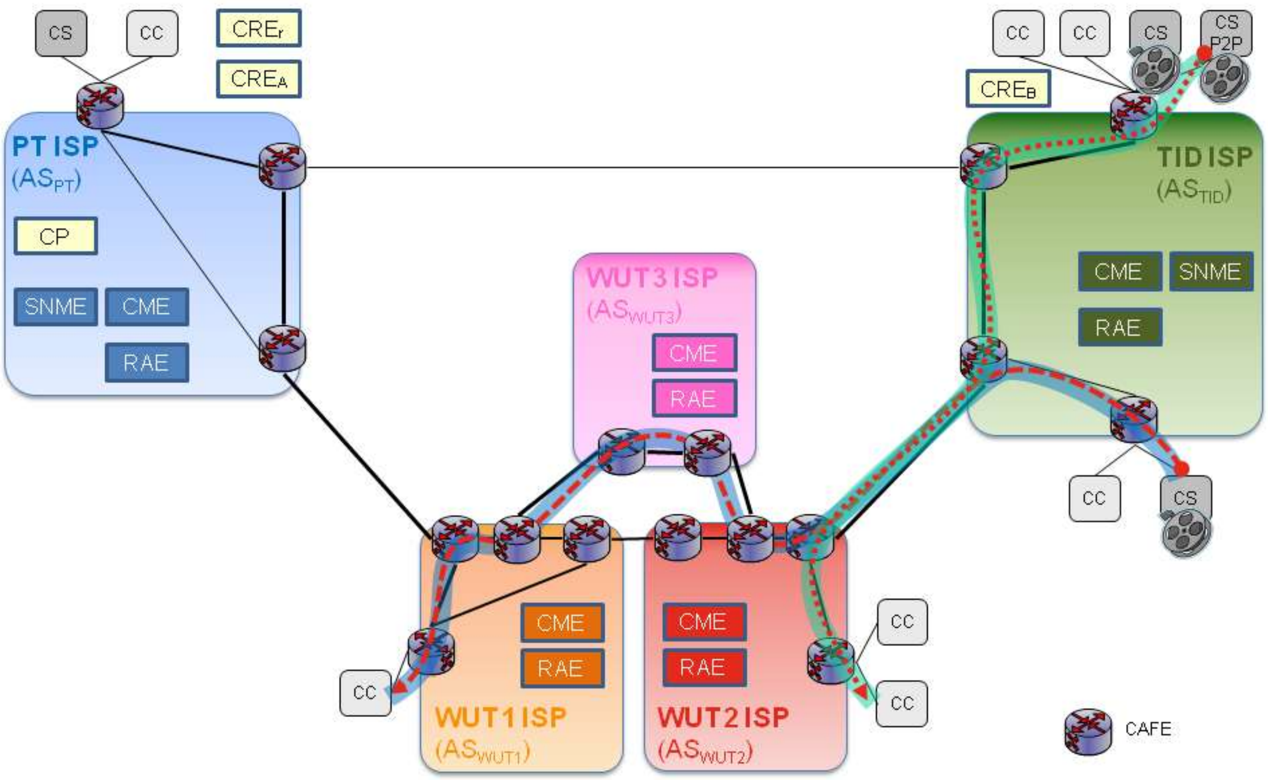

Figure 9: Path Established with a P2P CS

This description completes the basic operation of this Use Case. However, more complex scenarios can be tested taking this case as a starting point. The conditions and rules governing them are the same as those explained for Use Case 1 in section 4.1.3, so they will not be repeated here.

5 Test Plan

This section defines the tests to be carried out on the federated testbed described in Section 3. They are classified in two different categories:

- **Functional tests**. They will demonstrate the Use Cases described in Section 4.
- **Performance tests**. Individual tests carried out on the architectural elements in order to characterise them in terms of response time/maximum tolerated load/queries.

Scalability matters will be assessed in WP5, so they will not be considered here.

For each test case, the following pieces of information will be defined:

- **Test Identifier – A unique identifier.**
- **Objective of the test -** What the test intends to check.
- **Prerequisites -** Conditions that have to be fulfilled before starting the test.
- **Procedure -** Steps to be followed in order to carry out the test.
- **Expected Result -** What is expected to obtain after a successful completion of the test.

5.1 Functional Tests

The purpose of functional tests is the demonstration of the proper working of the Use Cases.

As a general rule, at least the entities enumerated in the general description of the testbed (CREs, CMEs, RAEs, SNMEs, CAFEs) and the Use Cases have to be installed, configured and activated.

5.1.1 Adaptable and efficient content distribution (Use Case 1)

This table gathers the test to be carried out for Use Case 1

Table 1: Tests for Use Case 1

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-001	How RAEs gather information about the paths linking the different ISPs (Routing Awareness).	Proper provisioning of PT's, TID's and WUT's domains to assure availability of required paths.	RAEs of neighbor ISPs exchange network information to learn how to reach other ISPs	Paths are propagated and stored in the CME.
FUN-UC1-002	Demonstrate how a Content Owner can publish contents in COMET (Content Publication).	Content Name to identify the Content. A CS where the Content can be hosted.	The Content Owner publishes a Content Record containing a Content Source (describing the characteristics and requirements of only the streaming servers distributing the Content).	The Content Record stored in the CRE Database.

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-003	Content Request Test	Publication of the content in COMET (creation of the Content Record) [FUN-UC1-002]. CC must have the same CoS with the Content Source included in the created Content Record.	The user invokes the CC by indicating the Content Name.	CC sends a query to CME in order to consume a content.
FUN-UC1-004	Name Resolution Test	The Content Request has to be done previously [FUN-UC1-003].	Retrieval of the Content Record associated to the Content Name from the CME.	The CME receives the Content Record with the defined Content Source.
FUN-UC1-005	How to obtain the load of the servers in the source.	The Content Record has been received by CME [FUN-UC1-004].	The CME sends a query to each CME of the servers' ISPs. Each CME of the servers' ISPs searches for the server status in its Servers Load DB (SNME).	CME of client's ISP receives the servers status information of CME of servers' ISPs.
FUN-UC1-006	Checking of the Decision Process	Paths are stored at CME [FUN-UC1-001]. Server Awareness [FUN-UC1-005].	Once all information about Content sources, paths and server's status is collected, CME automatically processes it.	CME decides on the optimal CS (It matches CC's CoS and it is in a low state of occupation) and the optimal path to provide the content (QoS requirements for CS matches the QoS requirements for path according to decision algorithms).
FUN-UC1-007	Path Configuration test	Optimal path selected [FUN-UC1-006].	Client's CME contacts server's CME to exchange information about the selected path. Server's CME configures thje edge CAFE attached to the CS.	Edge CAFÉ attached to CS is configured with the path information to reach thev CC

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-008	Content Delivery test	An end to end tunnel is established [FUN-UC1-007]. A content request has been made [FUN-UC1-003].	CC receives from the CME the information that will be used to request the Content from the selected server. CC launches the appropriate application in agreement with the information received from the CME.	An application is launched to download and reproduce the Content.
FUN-UC1-009	Demonstrate what happens when all CSs in the system are in HIGH state in terms of load	All CSs are overloaded. DecisionMaker is configured in strict mode.	All the processes related to content consumption previously described take place again, but the information provided by CMEs of server's ISP shows that all servers are overloaded.	No CS can be found for serving the content. CC waits for expiration time and ends execution
FUN-UC1-010	A new server is added	A CS where the Content can be hosted.	The Content Record is updated with another available CS in the Content Source.	The Content Record in the CRE Database has been updated.
FUN-UC1-011	A new user asks for the Content and the new server added is taken into account.	Content is published [FUN-UC1-010. All CSs are overloaded except new one.	The CC asks for the Content Name. All the processes related to content consumption previously described take place again.	A path is established between new CS and the user, and the content is retrieved.

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-012	A Content Owner changes the type of CSs distribution of the content.	The live event is concluded.	Content Owner removes the Content Streaming Sources and eliminates them. Content Owner adds new content VoD Sources and the VoD servers are powered on. These actualizations are done in the Content Record for the same Content Name.	The Content Record is updated in the CRE Database: the Content Streaming Sources has been erased and replaced by the VoD ones (the content is retrieved with the same Content Name).
FUN-UC1-013	VoD Content consumption using the same Content Record .	Content Record is updated [FUN-UC1-012].	The CC asks for the Content Name. All the process related to content consumption previously described takes place again.	The tunnel is established with one of the new VoD CS and the Content is retrieved.

5.1.1.1 Complex Scenarios

These complex scenarios will analyze in further detail the routing/decision capabilities of COMET, according to the CoS/BW/QoS constraints imposed in Content Sources and Paths.

Table 2: Test for Complex Scenarios

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-014	How CCs in different domains can access to the Same Content Name	Use Case 1 either for Streaming or VoD. A content is published consisting of one source and several servers CC at different ISPs	Clients at different ISPs request the same Content Name to their respective CMEs	The CS assigned by the CME to the CC belongs to the set defined in Content Record regardless the ISP the CCs are deployed

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-015	How the CME is capable of distinguishing users of different CoS and assigning them to different servers for Content retrieval.	Use Case 1 either for Streaming or VoD. A content is published consisting of two sources, with different CoS, A and B (i.e PR and BTBE). Clients with different CoS The population of CSs is different for each Content Source.	Clients with different CoS the same Content Name	The CS assigned by the CME to the CC belongs to the set defined for the User's CoS regardless of the ISP the CCs are deployed
FUN-UC1-016	How the CME manages the CoS rules.	Use Case 1 either for Streaming or VoD. A content is published consisting of three Content Sources, one for each CoS (PR, BTBE, BE). Clients for each of the three CoS. The population of CSs is different for each Content Source.	Clients with different CoS request the same Content Name to their CME.	Clients with lower CoS are not assigned Content Sources with higher CoS. Clients with higher CoS can be directed to Servers with lower CoS, if the higher ones are in HIGH state in terms of load.

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-017	How COMET can assign optimal paths according to the QoS/BW specified in the Content Record.	Use Case 1 either for Streaming or VoD. Published Content Record specifying QoS/BW requirements. A CC (for the CoS defined in the Content Record) and a CS defined in the Content Source. Several paths between the CC and CS provisioned with different QoS/BW. At least one matching the requirements in the Content Source.	The Client requests the Content Name to CME.	The CME answers with the Server defined in the Content Record. The Path assigned is the one provisioned to match the QoS/BW in the Content Source.
FUN-UC1-018	How COMET will not assign paths whose quality is lower than those defined in the Content Record.	Use Case 1 either for Streaming or VoD. Published available Content Record specifying QoS/BW requirements. A CC (for the CoS defined in the Content Record) and a CS defined in the Content Source. Several paths between the CC and CS provisioned with different QoS/BW. None matches the QoS/BW defined in the content Record.	The Client requests the Content Name to CME.	No response is returned to CC.

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-019	How, with every other network parameter the same, the path selection is ruled by servers's load status.	Use Case 1 either for Streaming or VoD A content Record specifying CoS/QoS/BW requirements. A CC (for the CoS defined in the Content Record) Several CS defined in the Content Source Several paths between the CC and CS provisioned with the same QoS/BW and mathing those in the CS.	The Client requests the Content Name to CME. One CS is in HIGH status in terms of load, the other not.	CME selects the server in lower status.

5.1.1.2 Point to multipoint streaming

The following table summarizes the test to be carried out in the case of point to multipoint subcase.

Table 3: Point to multipoint streaming Use Case

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC1-020	How Content Streaming Relay support point to multipoint streaming.	Publication of content suitable for point to multipoint streaming. Configuration of CSR to handle only 2 requests.	Publish content suitable for point to multipoint streaming. The CC in WUT2 domain asks for content available in point to multipoint streaming. Check address of content sever for first CC. New request for the same content arrives from second CC. Check address of content sever for second CC. New request for the same content arrives from third CC. Check address of content sever.	The first and the second request should be served by the CSR The third request should be served by CS.
FUN-UC1-021	Demonstrate ability to reduce bandwidth resource consumption.	Publication of content suitable for point to multipoint streaming. Configuration of CSR in domain WUT2 to handle a number of requests.	Publish content suitable for point to multipoint streaming. Generate a number of requests for the same content. Measure bandwidth consumption on inter-domain links in case when CSR is used and when it is not used.	The bandwidth consumption should be reduced proportionally to the number of request served by the CSR.

5.1.2 P2P Offloading (Use Case 4)

The table below summarizes the tests for the Use Case 4.

Table 4: Tests for Use Case 4

ID	Objective	Prerequisites	Procedure	Expected Results
FUN-UC4-001	How RAEs gather info about the paths linking the different ISPs (Routing Awareness).	Proper provisioning of PT's, TID's and WUT's domains to assure availability of required paths.	RAEs of neighbor ISPs exchange network information to know how to reach another ISPs.	Paths are propagated and stored in the CME database.
FUN-UC4-002	How a Content Owner distributes a Content publishing a Content Record with the number of Content Sources describing the characteristics and requirements of the streaming Servers.	Content Name to identify the Content. A CSs where the Content can be hosted.	The Content Owner publishes a Content Record with two Content Sources (Streaming & P2P ones) with different priorities.	A single content Record containing two Content Sources (Streaming & P2P) with info related to both types of CSs and their connection parameters will be published in the CRE.
FUN-UC4-003	How an end user requests the Content and retrieves the Content from the most appropriate CS (Streaming Distribution).	Path are stored at CME [FUN-UC4-001]. Content Publication [FUN-UC4-002].	The CC writes the Content Name assigned to the Content. All the process related to content consumption previously described in Use Case 1 takes place again.	The Content is retrieved through a streaming source (higher priority).
FUN-UC4-004	Show what happens when load is high in streaming servers.	Server awareness (VoD Servers are forced to a HIGH state).	end user requests the content using the same Content Name as in [FUN-UC4-003]. All the process related to content consumption previously described in Use Case 1 takes place again, but the information provided by CMEs of server's ISP shows that all streaming servers are overloaded.	The same result as in FUN-UC4-002 but in this case the Content is retrieved through a P2P Server (low quality sources).

5.2 Performance Tests

This section compiles the tests that will gather performance metrics for each of the functional entities of the Decoupled Approach, as defined in D3.2[4] and D4.2[5]

5.2.1 CRE Tests

This section gathers the performance tests related to the CRE

Table 5: Performance tests for CRE

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CRE-001	Evaluate capabilities of an authoritative CRE in terms of number of CRs than can be stored in its database.	Install 1 authoritative CRE. Prepare a stress tester that publishes content records to CRE, using the CP API, aiming to create Y CRs. CRs should be prepared to simulate a long tail distributrion (i.e 10% of CRs storing information of 100 CS vs. 90% of the CRs storing the information of 10 CS)	Run authoritative CRE, as well as the stress tester. Configure the stress tester to publish unlimited number of content records. Monitor when authoritative CRE stops creating and storing more content records and measure maximum number of content records.	Estimation of the maximum number of CRs that can be stored in one single authoritative CRE. <u>Target Value:</u> Billions (10^9) CRs in the overall CRE hierarchy.
PER-CRE-002	Assessment of authoritative CRE performance in terms of retrieval time according to the number of CRs stored in its database for sporadic queries.	Install 1 authoritative CRE. Prepare a CP application that publishes a specific number of content records to CRE, using the CP API. Based on the results of PER-PUB-001, configure the CP application for low occupation (20%), medium (70%) and high (90) emulating a long tail distribution. Implement a stress tester for requesting a CR.	Run the authoritative CRE, as well as the CP application. Prepare the CRE for low/ medium/ high load. Retrieve CRs for each load level and measure the retrieval time. Queries must be directly sent to the authoritative CRE as though the CME already knew content location.	Characterisation of retrieval times (minimum/ mean/ maximum and 95% percentile) vs CRE occupation

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CRE-003	Assessment of authoritative CRE performance in terms of retrieval time according to the number of queries per second and the number of CRs stored in its database.	1 authoritative CRE. Prepare a CP application that publishes a specific number of content records to CRE, using the CP API. Based on the results of PER-PUB-001, configure the CP application for low occupation (20%), medium (70%) and high (90). Implement a stress tester to launch N queries, in a Poisson distribution time pattern.	Run the authoritative CRE, as well as the CP application. Prepare the CRE for low/ medium/ high occupation. Launch sequences of N queries separated in a Poisson distribution time pattern. Measure the retrieval times for each query in the sequence. Queries must be directly sent to the authoritative CRE as though the CME already knew content location.	Characterisation of retrieval times (minimum/ mean/ maximum and 95 percentile) vs CRE occupation and query rates according to the Poisson Distribution
PER-CRE-004	Assessment of root CRE performance in terms of retrieval time according to number of naming authority records stored in its database for sporadic queries.	Install 1 root CRE. Prepare a CP application that publishes a specific number of naming authority records to CRE, using the CP API. Based on the results of PER-PUB-001, configure the CP application for low occupation (20%), medium (70%) and high (90) emulating a long tail distribution. Implement a stress tester for requesting a naming authority record.	Run root CRE, as well as the CP application. Prepare the CRE for low/ medium/ high occupation. Retrieve naming authority records for each occupation level and measure the retrieval time.	Characterisation of retrieval times (minimum/ mean/ maximum and 95 percentile) vs CRE occupation

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CRE-005	Assessment of root CRE performance in terms of retrieval time according number of queries per second and number of naming authority records stored in its database.	Install 1 root CRE. Prepare a CP application that publishes a specific number of naming authority records to CRE, using the CP API. Based on the results of PER-PUB-001, configure the CP application for low occupation (20%), medium (70%) and high (90). Implement a stress tester to launch N queries in a Poisson distribution time pattern.	Run root CRE, as well as the CP application. Prepare the CRE for low/ medium/ high occupation. Launch sequences of N queries in a Poisson distribution time pattern. Measure the retrieval times for each query in the sequence.	Characterisation of retrieval times (minimum/ average/ maximum and 95 percentile) vs CRE occupation and query rates according to the Poisson Distribution.
PER-CRE-006	Assesment of CP capabilities in terms of users that can simultaneously publish contents	Install 1 authoritative CRE and its respective CP. Implement a stress tester which emulates several end-users connected to CP.	Run both CRE and CP. Increase gradually number of end-users connected to CP and monitor when there is denial of service.	Maximum number of users connected to CP.
PER-CRE-007	Assesment of CP capabilities in terms of response times for publication and publication rates (pub/sec)	1 authoritative CRE. Prepare a CP application that publishes a specific number of content records to CRE, using the CP API. Based on the results of PER-PUB-001, configure the CP application for low occupation (20%), medium (70%) and high (90). Implement a tester to launch Publication Queries using a Poisson distribution time pattern	Run root CRE, as well as the CP application. Prepare the CRE for low/ medium/ high occupation. Launch sequences of N Publication queries in a Poisson distribution time pattern. Measure the retrieval times for each query in the sequence.	Characterisation of Publication times (minimum/ average/ maximum and 95 percentile) vs CRE occupation and publication rates according to the Poisson distribution.

5.2.2 CME Tests

This section gathers the performance tests related to the CME architectural entity. The will be focused on obtaining metrics for the different interfaces exposed by the CME:

- CC-CME interface
- Inter-CME interface
- RAE-CME interface

5.2.2.1 *Tests on CC-CME interface*

This section gathers the tests to be carried out on the CC-CME interface

Table 6: Performance tests for CC-CME interface

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CME-001	Evaluate CME performance in CC-CME interface	Deploy one CME. Implement a stress tester which emulates several CCs connected to the installed CME.	Run CME. Run multiple CCs querying multiple contents (no need to exist). Increase gradually the number of CCs sending requests to CME. Monitor success of requests, CPU and memory behavior, as well as response time.	Response Time of the CME (Minimum/Mean/Maximum and 95 percentile) Maximum number of requests per second that can be handled by CME in CC-CME interface.

5.2.2.2 *Tests on inter-CME interface (server awareness and path configuration)*

This section gathers the test to be carried out on the inter-CME interface

Table 7: Performance tests for interCME interface

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CME-002	Evaluate CME performance in inter-CME interface	Deploy one CME. Implement a stress tester which emulates multiple dummy client CME connected to the installed CME.	Run CME. Run multiple dummy client CME, sending dummy requests on the inter-CME interface. Increase gradually the number of dummy servers and monitor success of requests, response time, CPU and memory behavior.	Response Time of the CME (Minimum/Mean/ Maximum and 95 percentile) Maximum number of requests per second that can be handled by CME in inter-CME interface.

5.2.2.3 *Test on RAE-CME interface*

This section gathers the test to be carried out on the RAE-CME interface

Table 8: Performance tests for CME-RAE interface

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CME-003	Evaluate CME performance in RAE-CME interface	Deploy one CME. Implement a stress tester which emulates multiple requests from RAE connected to the installed CME.	Run CME. Run stress tester sending dummy updates on the RAE-CME interface. Increase gradually the rate of updates and monitor success of requests, CPU and memory behavior as well as response time	Maximum number of requests per second that can be handled by CME in RAE-CME interface.

[redacted]

[redacted]

- [redacted]
- [redacted]

[redacted]

[redacted]

[redacted]

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

[redacted]

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted] [redacted] [redacted]	[redacted] [redacted]	[redacted]

5.2.3 RAE Tests

This section gathers the performance tests to be carried out on the RAE

Table 11: Performance tests for RAE

ID	Objective	Prerequisites	Procedure	Expected Results

ID	Objective	Prerequisites	Procedure	Expected Results
PER-RAE-001	Check capabilities of RAE to handle 300 000 routing prefixes.	Install 2 RAEs on standalone hosts Prepare RAE configuration files with 1, 100, 10 000 and 300 000 prefixes (about 150 000 prefixes for each RAE).	Run RAEs with appropriate configuration files Check correctness of information provided by RAE via CME-RAE interface. Measure the routing convergence time for update and withdraw of prefixes. Moniotor usage of the CPU and memory.	RAEs should update the Known Routes Table (KRT) and Preferred Routes Table (PRT) with all prefixes. Information about paths should be provided via CME-RAE interface. The RAE should support up to 300 000 prefixes. The routing convergence time should be should be on the order of BGP-4 (expected few minutes)
PER-RAE-002	Evaluate routing convergence time after adding or removing of prefixes in a stub domain.	Prepare network with 3 core domains connected as full-mesh and one stub domain.	Run RAEs in each core domain Wait until routing is stable and check information about paths provided to CME. Start RAE in stub domain and measure routing convergence time after prefix advertisement. Next stop RAE in stub domain to measure routing convergence time after prefix withdrawal. (repeat procedure for different number of prefixes available in the core network)	The convergence time should be on the order of BGP-4 convergence time (expected few minutes)

ID	Objective	Prerequisites	Procedure	Expected Results
PER-RAE-003	Evaluate routing convergence time after inter-domain link failure between core domains	Prepare network with 3 core domains connected as full-mesh	Run RAEs in each core domain Wait until routing is stable and check information about paths provided to CME. Disconnect the inter-domain link between core domains and measure routing convergence time. Next recover the inter-domain link and measure routing convergence time. (repeat procedure for different number of prefixes available in the core network)	The convergence time should be on the order of BGP-4 convergence time (expected few minutes).

5.2.4 CAFE Tests

This section summarises the performance test to be carried out on CAFEs. Two different group of tests are analyzed, those for Edge CAFEs and those for core CAFEs, as well the tests for the CME-CAFE interface

5.2.4.1 Tests for Edge CAFE

This section gather the tests for edge CAFEs

Table 12: Performance tests for Edge CAFEs

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CAFE-001	To evaluate the forwarding throughput of Edge CAFE.	Install edge CAFE on standalone host with at least two 1 GbE interfaces. Prepare traffic generator/analyzer, e.g. Spirent TestCenter, to send IPv6 packets and receive and analyze COMET packets transferred through edge CAFE.	Configure edge CAFE to intercept IPv6 packets, encapsulate them with COMET header and forward to output interface. Increase the offered load until observing packet losses in the received packet stream. The throughput is defined as the maximum offered traffic without packet losses (RFC 2544). Measure the CPU load during the tests.	The edge CAFE should transfer packets with 1 Gbps throughput. Note, that obtained results should be compared with results for IP forwarding running on the same host.
PER-CAFE-002	To evaluate the performance of Edge CAFEfor increasing number of running flows	Install edge CAFE on standalone host with at least two 1 GbE interfaces. Prepare traffic generator/analyzer, e.g. Spirent TestCenter with CM-1G-D4 card, to generate a number of IPv6 flows and receive and analyze COMET packets transferred through edge CAFE.	Configure edge CAFE to intercept a number of IPv6 flows, classify them, encapsulate it with COMET headers and forward to output interface. Measure throughput under increasing number of running flows (up to 10 000). The throughput is defined as the maximum offered traffic without packet loss (RFC 2544) Measure the CPU load during the test.	The edge CAFE should be able to serve simultaneously at least 10,000 flows without packet losses

5.2.4.2 Tests for core CAFE

This section gathers the tests for core CAFEs

Table 13: Performance tests for core CAFEs

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CAFE-003	To evaluate the forwarding throughput of core CAFE	Install core CAFE on standalone host with at least tree 1 GbE interfaces and prepare configuration file. Prepare traffic generator/analyzer, e.g. Spirent TestCenter, to send COMET packets using different interfaces, and receive and analyze them after forwarding by core CAFE.	Run core CAFE using configuration file. Increase the offered load until observing packet losses in the received packet stream. The throughput is defined as the maximum offered traffic without packet losses (RFC 2544) Measure the CPU load during the tests.	The core CAFE should transfer packets with 1 Gbps throughput Note, that obtained results should be compared with results for IP forwarding running on the same host
PER-CAFE-004	To evaluate the performance of core CAFE for increasing number of running flows	Install core CAFE on standalone host with at least tree 1 GbE interfaces and prepare configuration file. Prepare traffic generator/analyzer, e.g. Spirent TestCenter with CM-1G-D4 card, to generate a number of COMET streams, receive and analyze COMET packets transferred through core CAFE.	Run CAFE using configuration file. Measure throughput under increasing number of running flows (up to 10 000). The throughput is defined as the maximum offered traffic without packet loss (RFC 2544) Measure the CPU load during the test.	The performance of core CAFE should be independent of the number of running flows

5.2.4.3 Tests on CME-CAFE interface

This section gathers the performance tests on the CME-CAFE interface

Table 14: Performance tests for CME-CAFE interface

ID	Objective	Prerequisites	Procedure	Expected Results
PER-CAFE-005	To assess time required for configuration of edge CAFE	Install edge CAFE on standalone host with at least one 1 GbE interface Develop tool which measures time required for configuration of edge CAFE.	Run edge CAFE. Measure time required for configuration of edge CAFE for at least 10000 flows. The measured values should include: min, mean and 95-percentile.	The configuration time should be in the order of a fraction of a second.

5.2.5 Global Performance Tests

This section summarises the performance test to be carried ou in order to characterize COMET as a whole from the point of view of performance.

Table 15: Global Performance tests

ID	Objective	Prerequisites	Procedure	Expected Results
PER-GLO-001a	Evaluate content resolution time (CRT) and success ratio (CRSR) for sporadic queries (queries are sent in intervals bigger than expected tolerable values for CRSR) when CCs and CSs are in the same ISP	Deploy and configure the entities required for content resolution (at least CRE, CME, SNME), for a single ISP Populate the CRE with CRs for multiple content names and for different occupation levels: low (20%) / medium (70%) high(90%). Populate the path tables in the CME with the path for accessing a CS inside the ISP. Populate the SNME database with CSs status. Implement a stress tester that simulates several CCs (Premium CoS) connected to a CME and can launch sequences of N queries in a Poisson distribution time pattern.	Run all entities with the required order. Launch resolution queries for different Content Names by using the CC simulator . Repeat the measurements for different CRE occupations. Measure CRT for each content request. Measure CRSR for content resolution.	Minimum/ Mean/ Maximum value of CRT and 95 percentile of CRT. Value of CRSR <u>Target values:</u> 95% percentile of CRT must not exceed 2,5 seconds. CRSR is 99.9%

ID	Objective	Prerequisites	Procedure	Expected Results
PER-GLO-001b	Evaluate content resolution response time (CRT) and success ratio (CRSR) for different queries rates (request per s) when CCs and CSs are in the same ISP	Same as PER-CME-001a	Run all entities with the required order. Launch resolution queries by using the CC simulator for sequences with different N in a Poisson distribution time pattern. Repeat the measurements for different CRE occupations. Measure CRT for each content request. Measure CRSR	Minimum /Mean /Max value of CRT and 95 percentile for different query rates Value of CRSR for different query rates <u>Target values:</u> 95% percentile of CRT must not exceed 2,5 seconds. CRSR must be 99.9%
PER-GLO-002	Evaluate content resolution response time (CRT) and success ratio (CRSR) for different query rates when CCs and CSs are in different domains	Deploy and configure the entities required for content resolution (at least CRE, CME, SNME), for a client and server ISP Populate the CRE with CRs for multiple content names and for different occupation levels: low (20%) / medium (70%) high(90%). Populate the path tables in the CME for 1 path between the ISPs Populate the SNME database with CSs status. Implement a stress tester that simulates several CCs (Premium CoS) connected to a CME and can launch sequences of N queries, in a Poisson distribution time pattern.	Run all entities in the required order. Launch resolution queries by using the CC simulator for sequences with different N in a Poisson distribution time pattern Repeat the measurements for different CRE occupations Repeat the measurements for ISP N hops away Measure CRT for each content request. Measure CRSR for content resolution	Minimum /Mean /Max value of CRT and 95 percentile for different query rates Value of CRSR for different query rates <u>Target values:</u> 95% percentile of CRT must not exceed 2,5 seconds. CRSR must be 99.9%

ID	Objective	Prerequisites	Procedure	Expected Results
PER-GLO-003	Evaluate content resolution response time CRT) for different query rates and success ratio (CRSR) when CCs and CSs are in different ISPs and the number of path increases	Deploy and configure the entities required for content resolution (at least CRE, CME, SNME), for a client and server ISP. Populate the CRE with CRs for multiple content names and for different occupation levels: low (20%) / medium (70%) high (90%). Populate the path tables in the CME for M paths between the ISPs. Populate the SNME database with CSs status. Implement a stress tester that simulates several CCs (Premium CoS) connected to a CME and can launch sequences of N queries in a Poisson distribution time pattern.	Run all entities with the required order. Launch resolution queries by using the CC simulator for sequences with different N in a Poisson distribution time pattern Repeat the measurements for different CRE occupations (at least medium and high) Repeat the measurements for for increasing M paths Measure CRT for each content request. Measure CRSR	Minimum /Mean /Max value of CRT and 95 percentile for different query rates Value of CRSR for different query rates Target values: 95% percentile of CRT must not exceed 2,5 seconds. CRSR must be 99.9%

ID	Objective	Prerequisites	Procedure	Expected Results
PER-GLO-004	Evaluate content resolution response time (CRT) and success ratio (CRSR) for different query rates when CCs and CSs are in different domains	Deploy and configure the entities required for content resolution (at least CRE, CME, SNME), for a client and server ISP Populate the CRE with CRs for multiple content names storing increasing amounts (M) of CS and for different occupation levels: low (20%) / medium (70%) high(90%). Populate the path tables in the CME for 1 path between the ISPs. Populate the SNME database with CSs status Implement a stress tester that simulates several CCs (Premium CoS) connected to a CME and can launch sequences of N queries in a Poisson distribution time pattern	Run all entities with the required order. Launch resolution queries by using the CC simulator for sequences with different N in a Poisson distribution time pattern Repeat the measurements requesting CRs with increasing M CSs. Repeat the measurements for different CRE occupations (at least medium and high) Measure CRT for each content request. Measure CRSR.	Minimum /Mean /Max value of CRT and 95 percentile for different query rates Value of CRSR for different query rates <u>Target values:</u> 95% percentile of CRT must not exceed 2,5 seconds. CRSR must be 99.9%

6 Mapping of requirements/metrics to tests

The objective of this section is to show how the requirements/metrics defined in COMET have been addressed with specific test from the set defined in this deliverable. Two types of requirement are taken into consideration

- Qualitative, as defined in D2.2 [3]
- Performance/Scalability metrics, as defined in D5.1 [6]

6.1 Qualitative Requirements

The following Table maps the requirements defined in D2.2 [3] with functional and performance tests defined in this deliverable.

Table 16: Mapping of Qualitative Requirements to test cases

ID	Category	System requirement	Test ID
1	**Global**	Content as a primitive	FUN-UC1-003 FUN-UC2-003 FUN-UC1-014 to FUN-UC1-015 FUN-UC1-019 FUN-UC4-002
2	**Global**	Global content naming and addressing	PER-PUB-001 Partially Applicable in WP5
3	**Global**	Open for future evolution of the Internet	Not Applicable in WP6 (architectural requirement)
4	**Global**	Scalable to be deployed in the largest ISPs	Not Applicable (WP5)
5	**Global**	Involvement of all Internet users as Content Creators	FUN-UC1-002
6	**Global**	Graceful switching of the content delivery path without impact on the application-layer	Not covered by the selected Use Cases
7	**Content Consumer**	Access independent from content location	FUN-UC1-011 FUN-UC1-014 FUN-UC1-015
8	**Content Consumer**	Content ID independent from way distribution and nature of content	FUN-UC1-010 FUN-UC1-012 FUN-UC1-013 FUN-UC1-014 FUN-UC1-015 FUN-UC1-019 FUN-UC4-002 FUN-UC4-004

ID	Category	System requirement	Test ID
9	**Content Consumer**	User unawareness	FUN-UC1-003 FUN-UC1-008 FUN-UC1-013 FUN-UC1-020 FUN-UC4-003
10	**Content Consumer**	CC able to declare his capabilities	FUN-UC1-003 to FUN-UC1-008 FUN-UC1-015 to FUN-UC1-019 FUN-UC4-001 to FUN-UC4
11	**Content Consumer**	CC will obtain all necessary parameters to invoke the application level requests	FUN-UC1-008 FUN-UC1-013 FUN-UC4-003 FUN-UC1-020 FUN-UC4-003
12	**Content Provider**	Interface to update the content properties	FUN-UC1-010 FUN-UC1-012
13	**Content Provider**	Capability of establishing policies to enforce the way to deliver contents	FUN-UC1-002 FUN-UC1-010 FUN-UC1-012 FUN-UC1-015 to FUN-UC-017 FUN-UC4-002
14	**CMP**	Global content resolution architecture	FUN-UC1-004 FUN-UC1-014
15	**CMP**	Integrated traffic and resource management solution to increase network efficiency and content delivery	FUN-UC1-001 FUN-UC1-004 FUN-UC1-015 to FUN-UC1-19 FUN-UC4-001 FUN-UC4-004
16	**CMP**	Information gathering system	Not Implemented yet
17	**CMP**	Efficient protocol interfaces	FUN-UC1-002 FUN-UC1-003 FUN-UC1-008 FUN-UC1-015 to FUN-UC1-019 FUN-UC4-002 FUN-UC4-003
18	**CMP**	Capability of dynamically modify servers location information	FUN-UC1-010 FUN-UC1-012 FUN-UC1-015 to FUN-UC1-019 FUN-UC4-002
19	**CMP**	Possibility of registering different ways of distribution	FUN-UC1-010 FUN-UC1-012 FUN-UC1-015 to FUN-UC1-019 FUN-UC4-002

ID	Category	System requirement	Test ID
20	**CMP**	Network conditions and routing information awareness	Only for long term information FUN-UC1-001 FUN-UC4-001 PER-RAE-001 PER-RAE-002 PER-RAE-003
21	**CMP**	Interaction between the Content Mediation Servers and the Content Aware Forwarders to enforce content delivery	FUN-UC1-005 FUN-UC1-007 FUN-UC1-015 to FUN-UC1-019 FUN-UC4-003 PER-CAFE-005
22	**CMP**	CMP able to request the enforcement of QoS and multicast in the network	FUN-UC1-015 to FUN-UC1-019 FUN-UC1-019
23	**CFP**	Content forwarding architecture able to reach IP-based forwarding speeds	PER-CAFE-001 to PER-CAFE-004
24	**CFP**	Elements in CFP able to support QoS-aware content delivery	FUN-UC1-001 FUN-UC1-015 to FUN-UC1-019 FUN-UC1-001
25	**CFP**	Elements in CFP able to support point-to-multipoint content delivery	FUN-UC1-020 FUN-UC1-021
26	**CFP**	Content may be cached to optimize network resource usage	Not in decoupled Approach
27	**CFP**	Interaction between the CFP and the CMP to provide information on network conditions and, optionally, routing information	Only long term info FUN-UC1-001 FUN-UC1-001

6.2 Performance/Scalability Metrics

The following Table maps the Performance ScalabilityMetrics defined in D5.1 [6] as well as the functional and performance tests defined in this deliverable that cover them.

Table 17: Performance/Escalability Metrics

Type	Metric	Test ID
Content retrieval	Content Retrieval Latency (CRL) (expressed by 95 percentile of CRL)	PER-GLO-001 to PER GLO-004
	Content Resolution Time (CRT) (expressed by 95 percentile of CRT)	PER-GLO-001 to PER GLO-004
	Content retrieval success ratio (CRSR) (expressed by % of successful requests)	PER-GLO-001 to PER GLO-004
	Content resolution signaling overhead, expressed by number of traversed ASes	Evaluated in WP5
	Maximum request rate (expressed in [req/sec])	PER-GLO-001 to PER GLO-004

Type	Metric	Test ID
Content publication	Maximum number of content records stored in a single Authoritative CRE	PER-CRE-001
	Maximum number of users connected to the CP	PER-CRE-006
	Maximum publication rate (expressed in [pub/sec])	PER-CRE-007
Content delivery	Hop count the number of AS-level hops required for a particular content to reach the client from the server.	Evaluated in WP5
	Reduction of bandwidth consumption defines the bandwidth saving achieved under multicast over unicast. This is represented in relation to the number of hops.	Partially demonstrated in FUN-UC1-021 Evaluated in WP5
CRE performance	Maximum query rate of root CRE (expressed in [req/sec])	PER-CRE-005
	Response time of root CRE (expressed by 95 percentile)	PER-CRE-004, PER-CRE-005
	Maximum query rate of Authoritative CRE (expressed in [req/sec])	PER-CRE-002
	Response time of Authoritative CRE (expressed by 95 percentile)	PER-CRE-002, PER-CRE-003
SNME performance	[redacted]	PER-SNME-008, PER-SNME-009
	[redacted]	PER-SNME-008, PER-SNME-009
	[redacted]	PER-SNME-001 to 003
RAE performance	Maximum number of stored network prefixes	PER-RAE-001
	Routing convergence time	PER-RAE-002, PER-RAE-003
CAFE performance	Lossless throughput	PER-CAFE-001 PER-CAFE-003
	Number of simultaneous flows	PER-CAFE-002 PER-CAFE-004
	Edge CAFE configuration latency	PER-CAFE-005
	Size of Forwarding Information Base (FIB)	Evaluated in WP5
	Size of COMET header	Evaluated in WP5

Type	Metric	Test ID
CME performance	Maximum request rate of client CME (expressed in [req/sec])	PER-CME-001
	Response time of client CME (expressed by 95 percentile)	PER-CME-001
	Maximum request rate of server CME (expressed in [req/sec])	PER-CME-002
	Response time of server CME (expressed by 95 percentile)	PER-CME-002

7 COMET-ENVISION Interface

7.1 Introduction

In this section, we describe the plan for the joint COMET-ENVISION collaboration. While both projects focus on various aspects of digital data content in the Internet (content access, dissemination, delivery etc.), the high-level approaches are different. In the COMET project, we tackle the issue on content dissemination via an overlay approach at the network level resulting in a 2-plane approach aiming to mediate the delivery of Internet content via native COMET network entities. On the other hand, the ENVISION project deals with the problem by developing techniques for the content delivery at the application layer and by fostering the collaboration between the applications and the underlying ISP networks to achieve the co-optimisation of the often misaligned application and network performance objectives.

We identify a possible collaboration between the two projects by exploiting the fact that both COMET and ENVISION propose the creation of novel entities owned by ISPs which can then collaborate to gain better network awareness and thus, achieve better content delivery performance. We describe the scenario considered in the next sub-section and present our integration plan after that.

7.2 Integration Scenario

In ENVISION, an ALTO-compliant server, called the CINA server, is developed with the purpose of facilitating applications in finding better content delivery options. Each ENVISION-enabled domain deploys a CINA server through which, among other things, it allows the applications to query about network connections, annotated or ranked using different costs. These costs may reflect network performance properties (e.g. routing hop count, estimated delay, etc.) and/or ISP preferences (e.g. cost of transit links) and are used to influence the decisions at the application layer. In COMET, these ENVISION costs can be used as additional criterion in the selection of a content delivery path. For a particular content item, several paths may exist to reach a server with the content, and as the forwarding is performed over CAFE nodes the connection to be ranked by the CINA server is the connection to the next-hop CAFE rather than the connection to the final destination. The CINA server would therefore provide costs for the next-hop domain for all of the candidate paths.

In COMET, content delivery paths are mediated by the content mediation function (CMF) located at the content mediation entity (CME) which again, each is owned by the respective ISP. A highly extensible 2-phase multi-criteria decision algorithm (cf. [4]) has been developed in COMET to take into account various performance metrics when choosing the optimal content delivery path. This algorithm is implemented within the CME.

The setup of the integration scenario enables the CINA server to feed to its local CME the costs it maintains such that the decision algorithm within CME can exploit the added information when optimizing the content delivery to end users.

7.3 Integration Plan

7.3.1 Basic Setup

To enable collaboration between the two projects, we envisage the creation of a COMET-ENVISION interface that enables communication between the relevant entities from both sides. Specifically, this interface is foreseen to be between ENVISION's CINA server and COMET's CME. Via this interface, the CINA server can supply certain network information obtained within the ENVISION system to the COMET system, thus allowing the related COMET functions to perform their optimization in a more timely and informed manner.

We illustrate the integration setting in Figure 10 within the COMET federated testbed. While the figure shows all domains with end users to have a CINA server, in any particular evaluation scenario not all domains need to be ENVISION-enabled. For example, if some content item requested by CC@AS_{WUT1} can be found through AS_{PT}, AS_{WUT2} or AS_{WUT3}, the CINA server would return the corresponding costs of receiving incoming traffic through these domains to CME

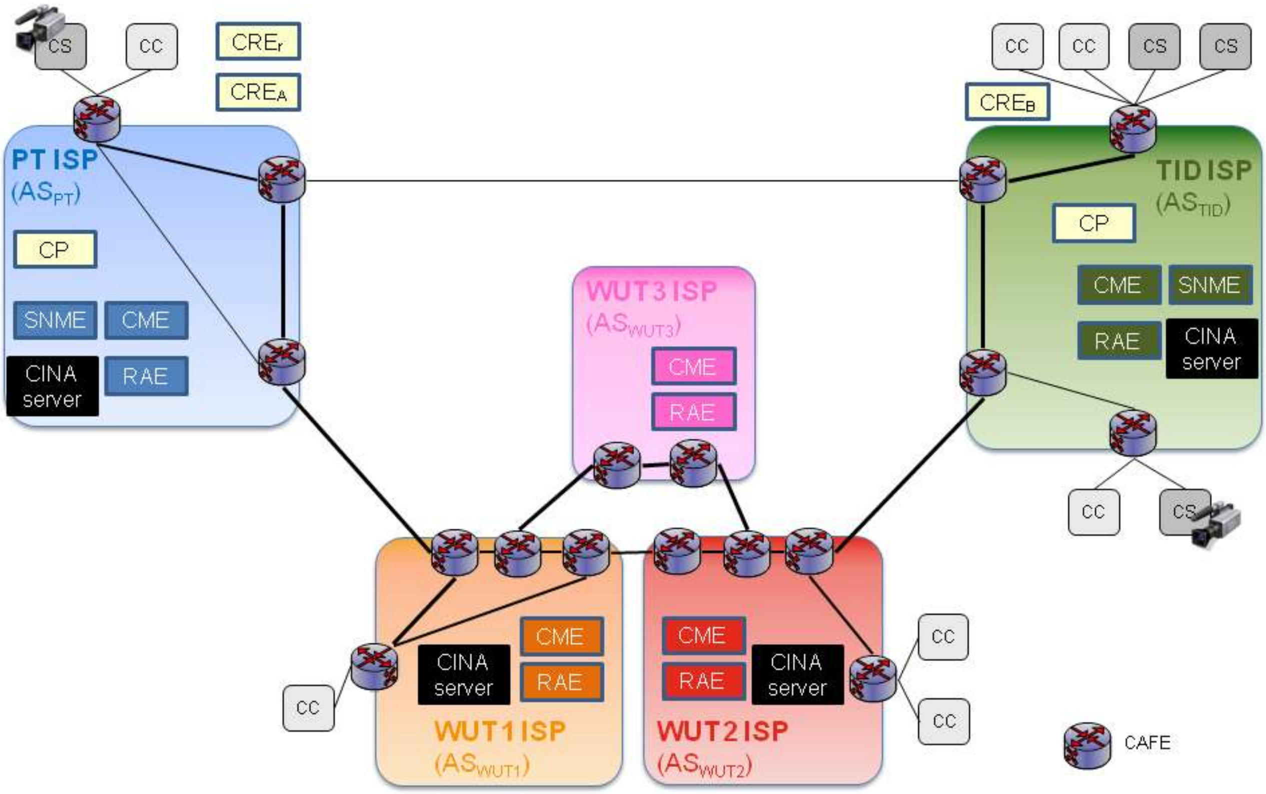

Figure 10: COMET-ENVISION integration

7.3.2 Comet-Envision Interface

The envisioned COMET-ENVISION interface for cost discovery follows the ENVISION CINA specifications. In the case of this integration scenario, we will use the ALTO-compatible endpoint cost service, see section 7.7.5 of ALTO protocol (http://tools.ietf.org/html/draft-ietf-alto-protocol-10). The endpoint in this case is the IP address of the CAFE that is the next-hop on a candidate path. The HTTP request and response would be for example:

```
POST /endpointcost/lookup HTTP/1.1
Host: cina.wuit1.com
Content-Length: [TODO]
Content-Type: application/alto-endpointcostparams+json
Accept:application/alto-endpointcost+json,application/alto-error+json
 {
   "cost-mode" : "ordinal",
   "cost-type" : "administrativecost",
   "endpoints" : {
```

```
    "dsts": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34",
      "ipv4:203.0.113.45"
    ]
  }
}

HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto-endpointcost+json
{
  "meta" : {},
  "data" : {
    "cost-mode" : "ordinal",
    "cost-type" : " administrativecost",
    "map" : {
      "ipv4:192.0.2.2": {
        "ipv4:192.0.2.89"    : 1,
        "ipv4:198.51.100.34" : 2,
        "ipv4:203.0.113.45"  : 3
      }
    }
  }
}
```

8 Summary and Conclusions

As it was stated in the introduction the aim of this deliverable is threefold.

First, to devise the guidelines that rule the setup of the envisaged COMET federated testbed, with the aim of creating a testing ground as close as possible to a real scenario, consisting of a multiplicity of ISPs and paths linking them, where the basic concepts and architectural elements of COMET are deployed and tested.

With this aim in mind, the envisioned COMET Testbed will incorporate five ISPs (three in WUT and one respectively in TID and PT) with the possibility of multiple paths between any two given locations, thus enabling paths' qualification with different BW/QoS parameters in order to test COMET's Path Discovery, Decision and Configuration Capabilities.

Besides, the testbed layout is flexible enough to set up a hierarchy of CRE, in order to make visible the decoupling between Content Resolution and Content Consumption, a feature of the decoupled approach. The testbed layout also allows deploying CCs and CSs in almost any location, such that complex content distribution and consumption scenarios could be devised and tested (i.e. a scenario with a multiplicity of servers distributing the same content, located in different ISPs and clients consuming the content from different location and ISPs) beyond those sketched in this document.

A second important task was the decision on which Use Cases from those defined in D2.1 [2] and D2.2 [3] are going to be tested and demonstrated on the federated testbed. The chosen ones have been the **Use Case 1: Adaptable and efficient content distribution** (with a subcase studying the multicast in a client's ISP) and **User Case 4: P2P Offloading,** on the basis that they cover most of the basic functionalities and requirements of the final COMET System. Additionally, those Use Cases have been analysed in detail in order to show how they can be demonstrated in the envisaged testbed and which architectural elements/capabilities will be tested in turn by each of them. As demonstration activities are performed over the federated testbed, new testing possibilities can and will surely arise, whose results will be reflected in D6.2.

As the third main contribution, a set of test cases have been defined, focusing in all the functional tests related to the Use Cases previously described. This ensures that every aspect of their functionality is assessed. At the same time, performance matters have also been taken into account by defining a set of tests for each architectural entity, so that they can be characterised in terms of response times and/or maximum tolerable amount of queries. Scalability considerations, though, are out of the scope of WP6, and will be analysed in WP5 and related deliverables.

In order to link the results in the deliverable with the work carried out in other work packages (mainly WP2, architecture and WP5 Validation and Evaluation), the qualitative defined in D2.2 [3] and the performance metrics defined in D5.1 [6] have been gathered in tables that show the functional and performance tests addressing them, in order to prove that COMET's intended functionality is fully covered by the testing and demonstration activities to be carried out in the scope of T6.2 and defined in this deliverable.

Finally, it has been sketched how COMET and ENVISION could be integrated. The solution consists in using the path-cost information gathered by ENVISION in order to refine the decision process used by COMET to assign a CS (and a path) to a CC for retrieving a specific content. This will involve the deployment of ENVISION's CINA module in the ISP set up in the the federated testbed, as well as developing an interface for communicating with the CMEs.

9 References

[1] The EU FP6 AGAVE Project, "A Lightweight Approach for Viable End-to-end IPBasedQoS Services", http://www.ist-agave.org

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

[3] COMET Deliverable D2.2 High-Level Architecture of the COMET System

[4] COMET Deliverable D3.2 Final Specification of Mechanisms, Protocols and Algorithms for the Content Mediation System

[5] COMET Deliverable D4.2 Final Specifications of the Mechanisms, Protocols and Algorithms for Enhanced Network Platforms

[6] COMET Deliverable D5.1 Integration of COMET Prototype and Adaptation of Applications

[7] http://www.videolan.org/

[8] http://www.utorrent.com/

[9] S. Bradner, J. McQuaid, Benchmarking Methodology for Network Interconnect Devices, IETF RFC 2544, March 1999

10 Abbreviations

AS	Autonomous System
BE	Best Effort
BtBE	Better than Best Effort
BW	Bandwidth
CAFE	Content-Aware Forwarding Entity
CINA	Collaboration Interface between Network and Application
CC	Content Client
CFP	Content Forwarding Plane
CME	Content Mediation Entity
CMP	Content Mediation Plane
COMET	COntent Mediator architecture for content-aware nETworks
CoS	Class of Service
CP	Content Publisher
CRE	Content Resolucion Entity
CRL	Content Retrieval Latency
CRSR	Content Retrieval Success Ratio
CRT	Content Resolution Time
CS	Content Server
CSR	Content Streaming Relay
DL	Direct Download
ENVISION	Co-optimisation of overlay applications and underlying networks
FIB	Forwarding Information Base
IP	Internet Protocol
ISP	Internet Service Provider
LAN	Local Area Network
Pr	Premium
PT	PrimeTel
QoS	Quality of Service
RAE	Route Awareness Entity
SIC	Server Information Collector
SLA	Service Level Agreement
SNME	Server and Network Management Element
STREP	Specific Targeted Research Project
TID	Telefonica I+D
URL	Universal Resource Locator
VLAN	Virtual LAN

VLC	VideoLan Player
VoD	Video On Demand
WUT	Warsaw University of Technology

11 Acknowledgements

This deliverable was made possible due to the large and open help of the WP5 team of the COMET team within this STREP . Many thanks to all of them.