COmtent Mediator architecture for content-aware nETworks

COMET, the Decoupled Approach: Mediating between Content Producers and Consumers

David Flórez (dflorez@tid.es)

Telefonica I+D

Rationale

Content Access still requires users' previous knowledge of machines hosting it.

- Need of searching mediators
- Contents "trapped" in isolated communities

Network is unaware of the Contents

- Carried to their final destination by using best efforts strategies
- Little or no use of QoS capabilities,
  - Especially if content and user are located at different ISPs

Proposal (1/2)

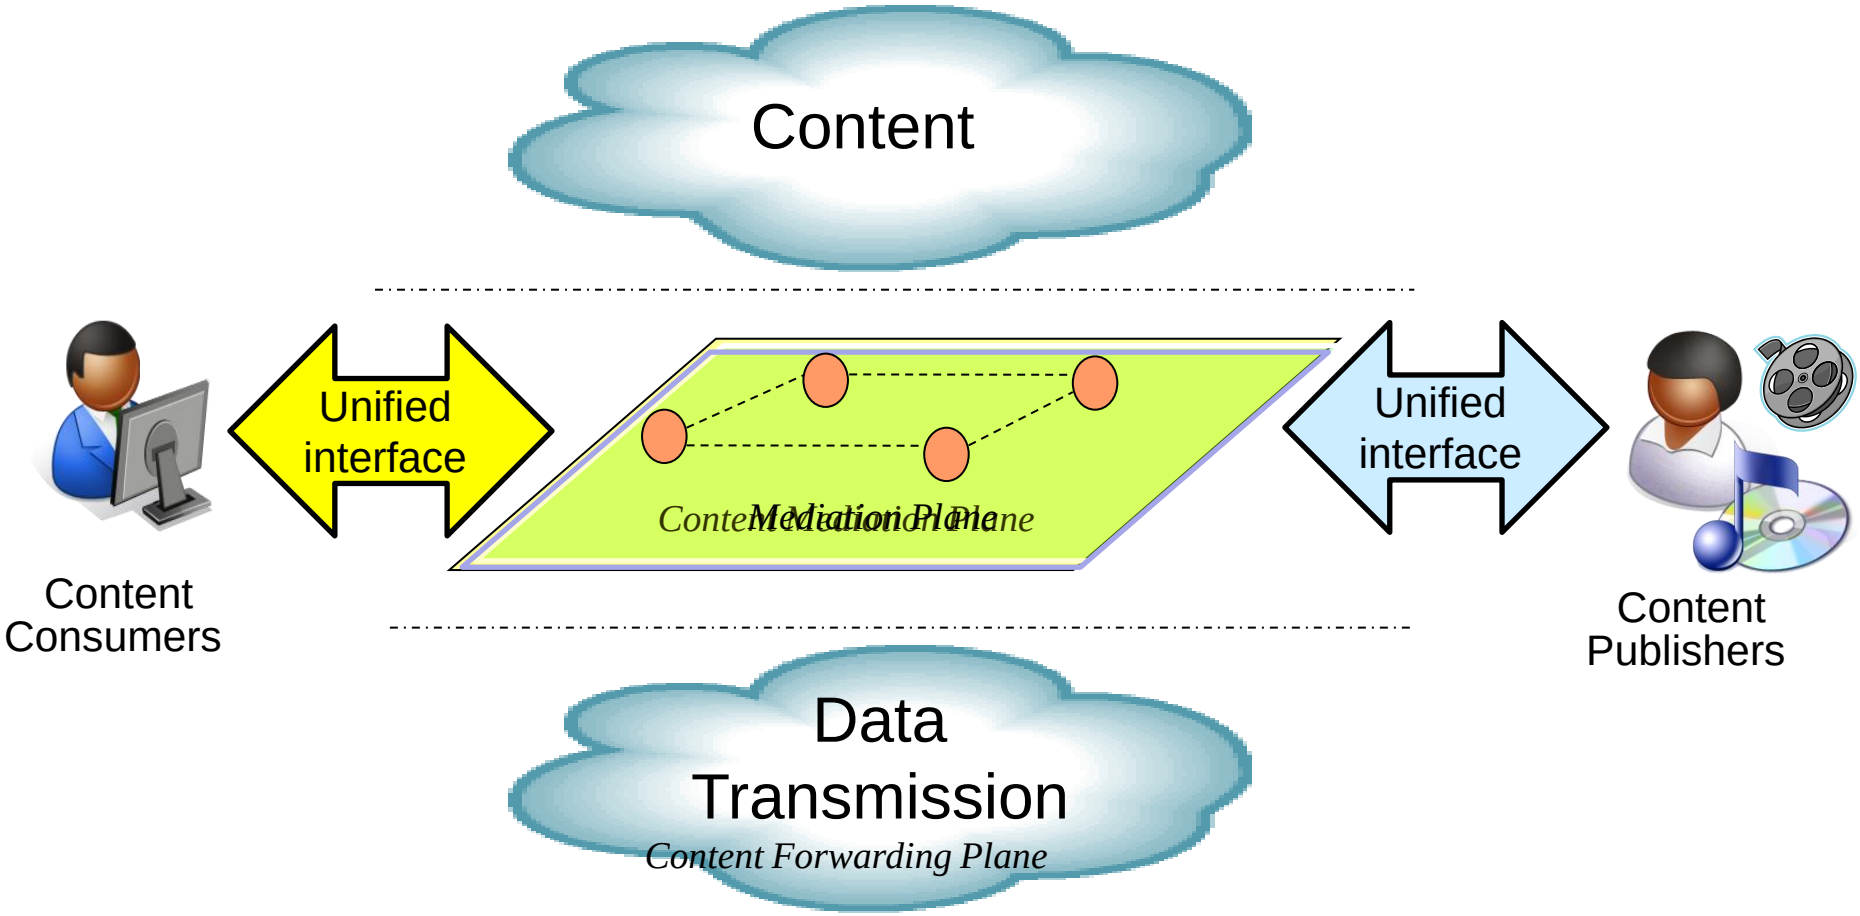

Proposal (2/2)

Mediation Plane, offering an unified interface

- For Content Providers to made the network aware of the contents
  - Assigning a unique identifier for the content, regardless where it is located and how it will be transmitted
- For Users, in order to request the retrieval of those contents by using the unique identifier
- For the Network,
  - Location of the Contents
  - Routing Information
  - Servers' status (load, transport protocols)

Content Forwarding Plane

- Sets up the channel between the user and the server hosting the requested contents
- Best Path in terms of QoS end to end, according to
  - Server's status
  - Content Requirements
  - Service Level subscribed by the user

The Decoupled Approach

- Content resolution and path setup are separate as in today's DNS-based Internet
- Allows graceful transition from today's host-centric Internet
- Content resolution is based on a global content directory system

General Architecture

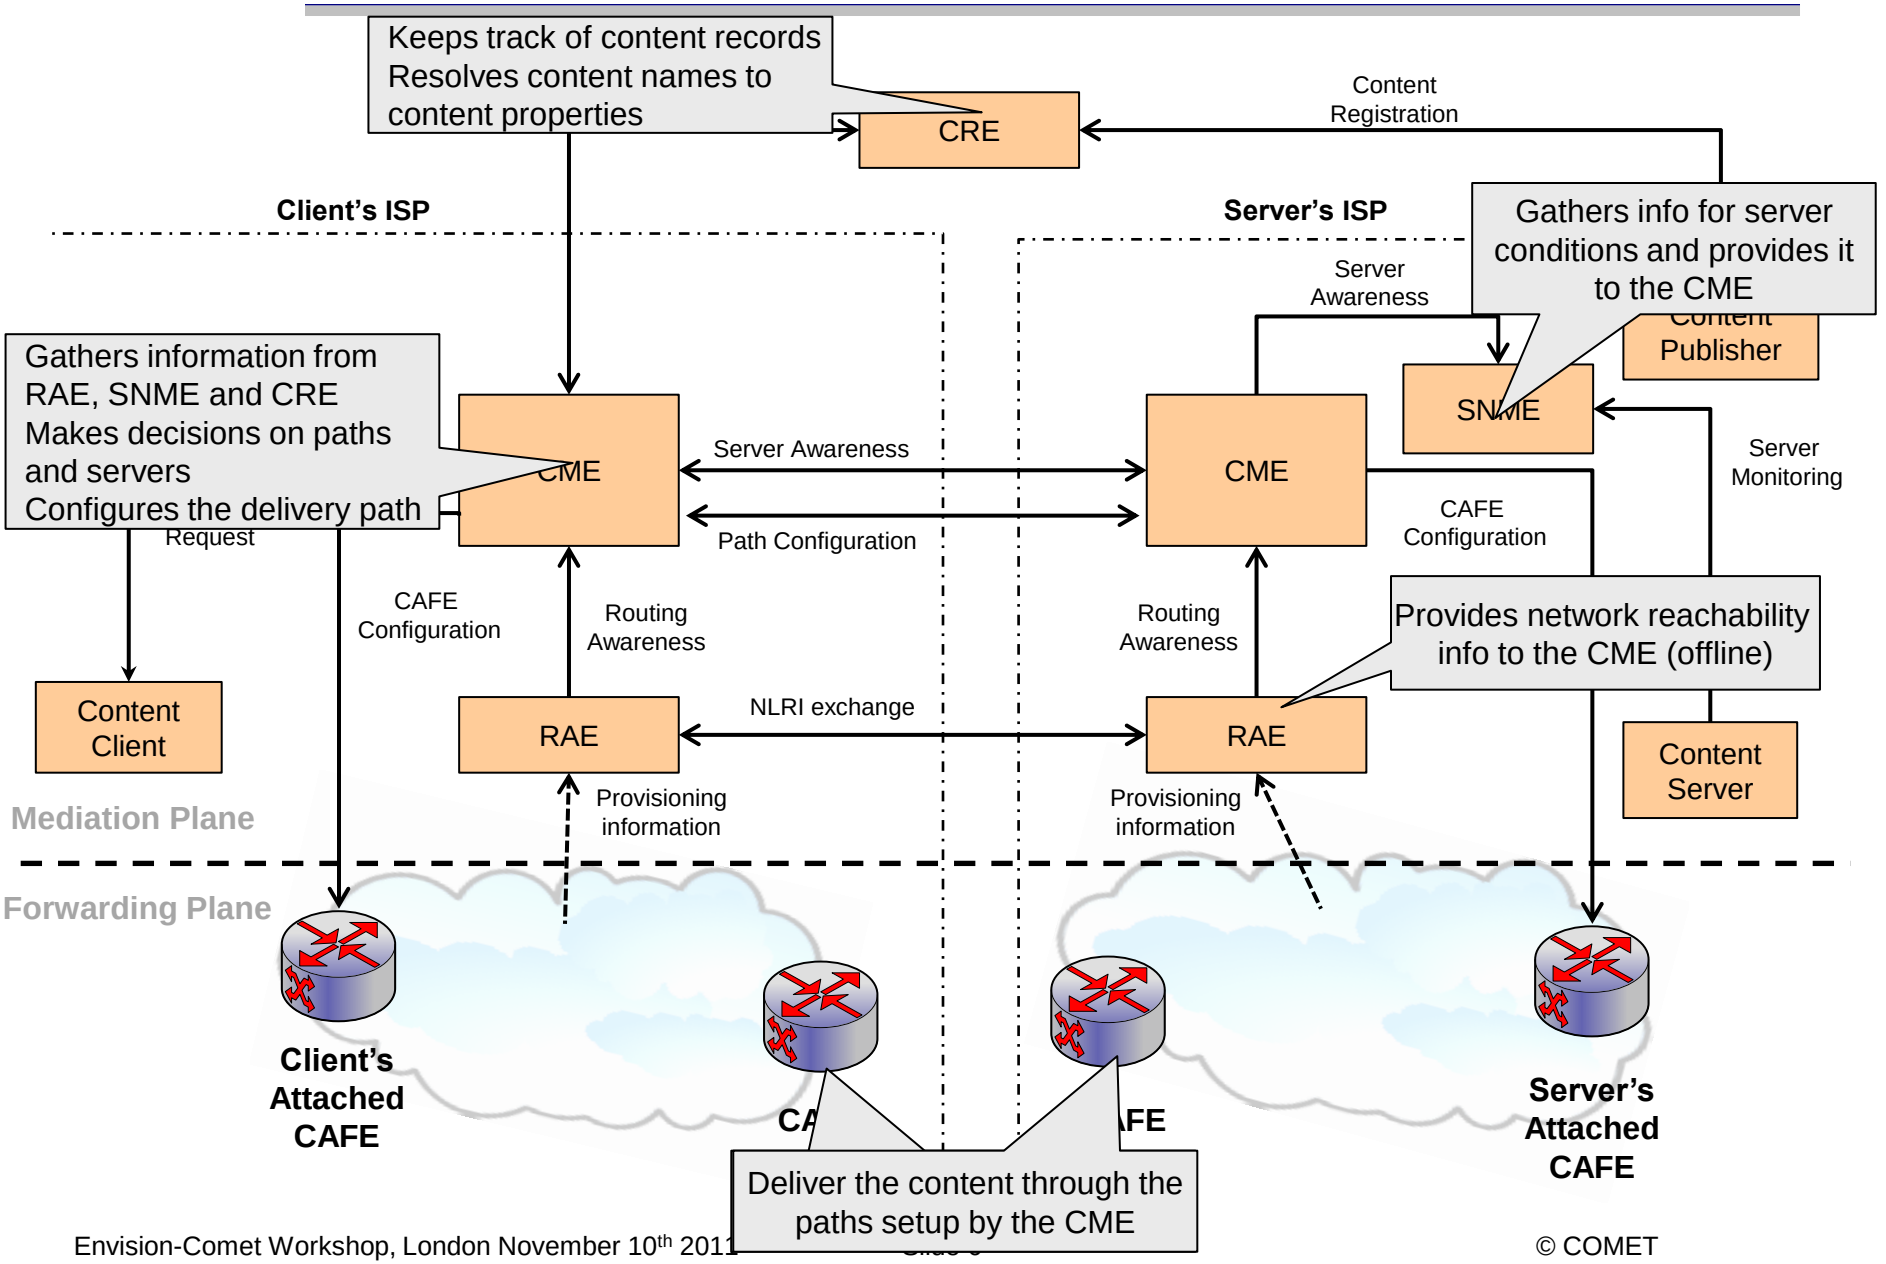

Remarks

- CME/SNME/RAE/CAFE will be located at each ISP
- CREs are “outside” the ISPs, hence the decoupled name.
  - A single CRE can manage several ISP
  - There will exist a hierarchy of CREs (DNS-like)
- There will be more than one CAFE per ISP
  - Located at the links with other ISPs
  - For controlling Content Clients and Content Servers
- There may be more than one SNME
  - Depending on the Content Servers’ density.

- *Server Awareness*
  - How the Content Server informs the SNME with their load status
- *Route Awareness*
  - How the RAE informs the CME about qualified available routes to hop from ISP to ISP till the intended destination
- *Content Publication*
  - How the Content Publisher makes contents available throughout Comet by publishing Content Records to CRE
- *Content Consumption*
  - How a Content Client will retrieve a content mediated by Comet

Server Awareness (1/2)

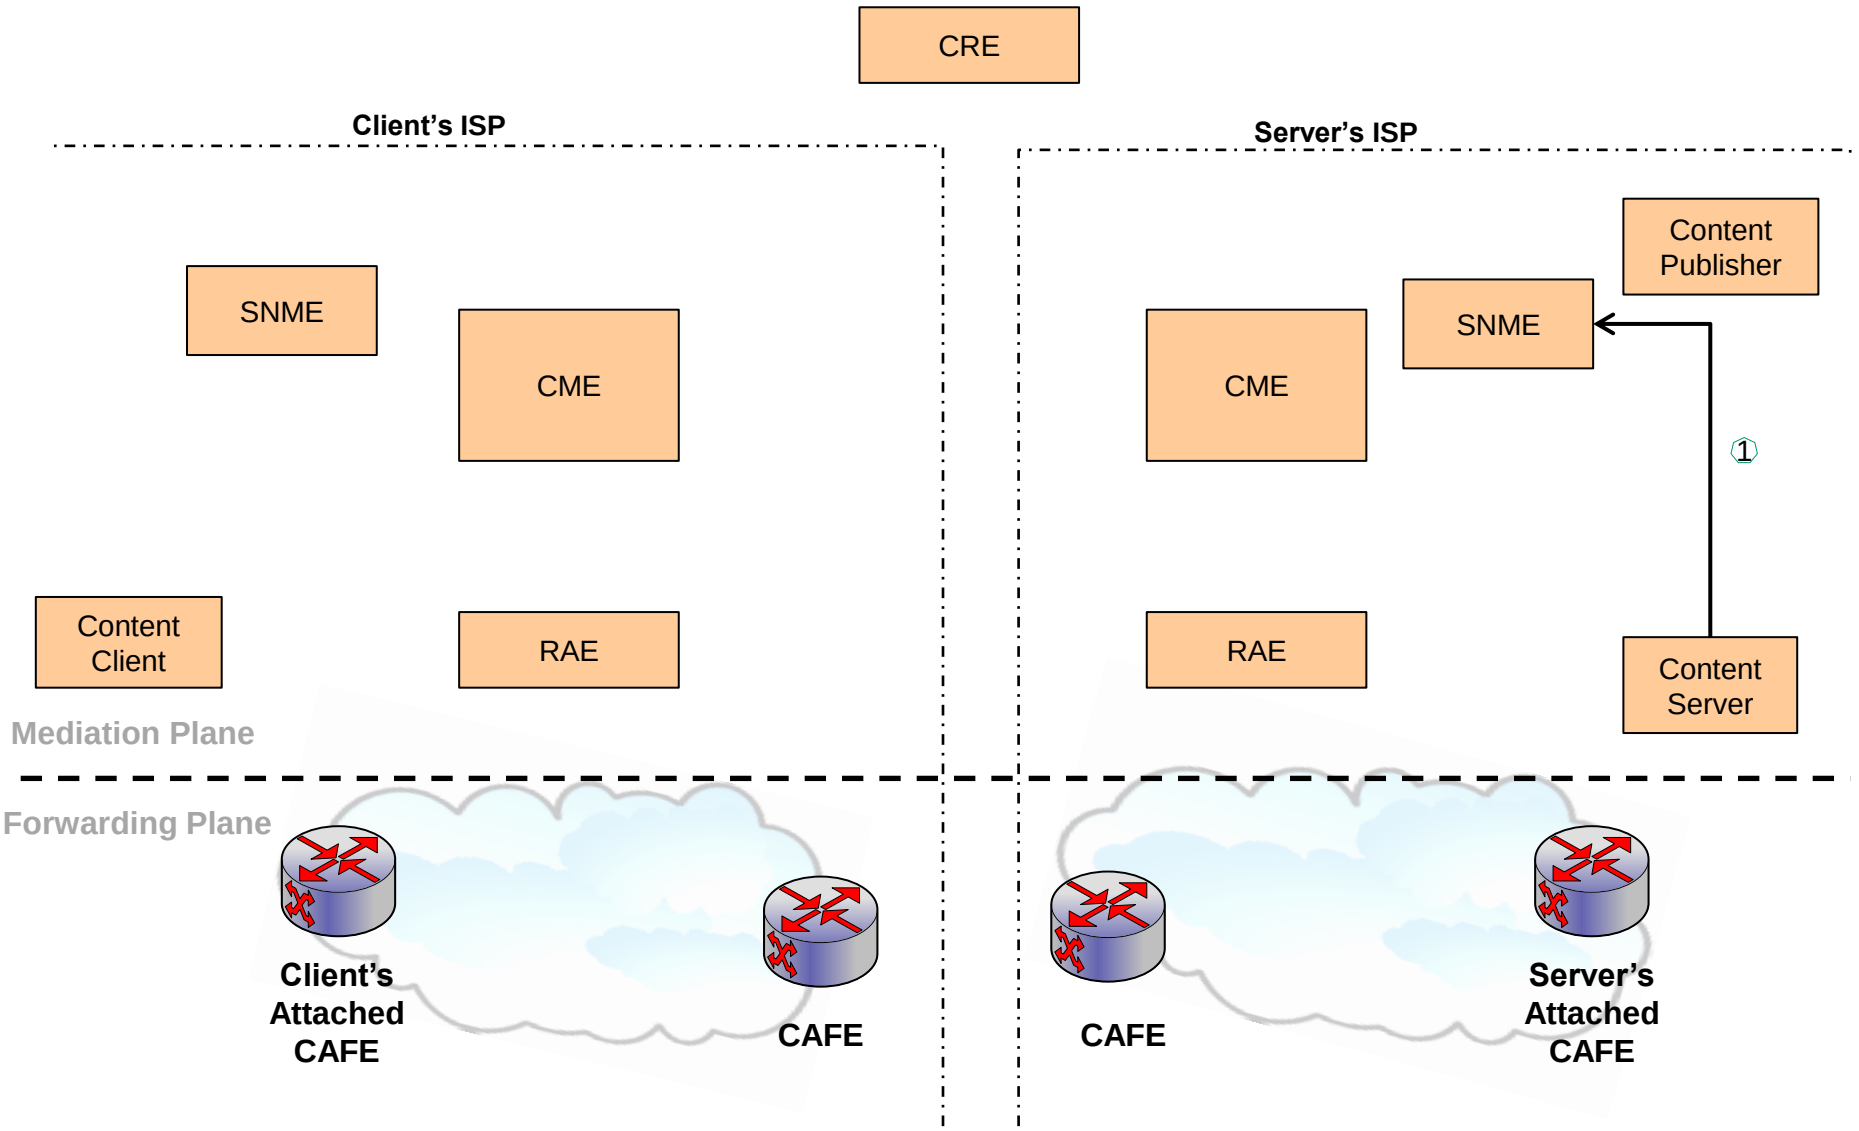

Server Awareness (2/2)

1. Content Server assesses its load and sends the result to SNME, where it is stored
   - A number of load parameters will be taken into account
     - Typically, CPU Load, Memory Allocation, Free Bandwidth
   - Status is coded into a predefined set of basic states
     - High/Medium/Low
       - High, meaning this server should be avoided for retrieval

This is an continual offline process. it is not triggered by events in the CME

Routing Awareness (1/2)

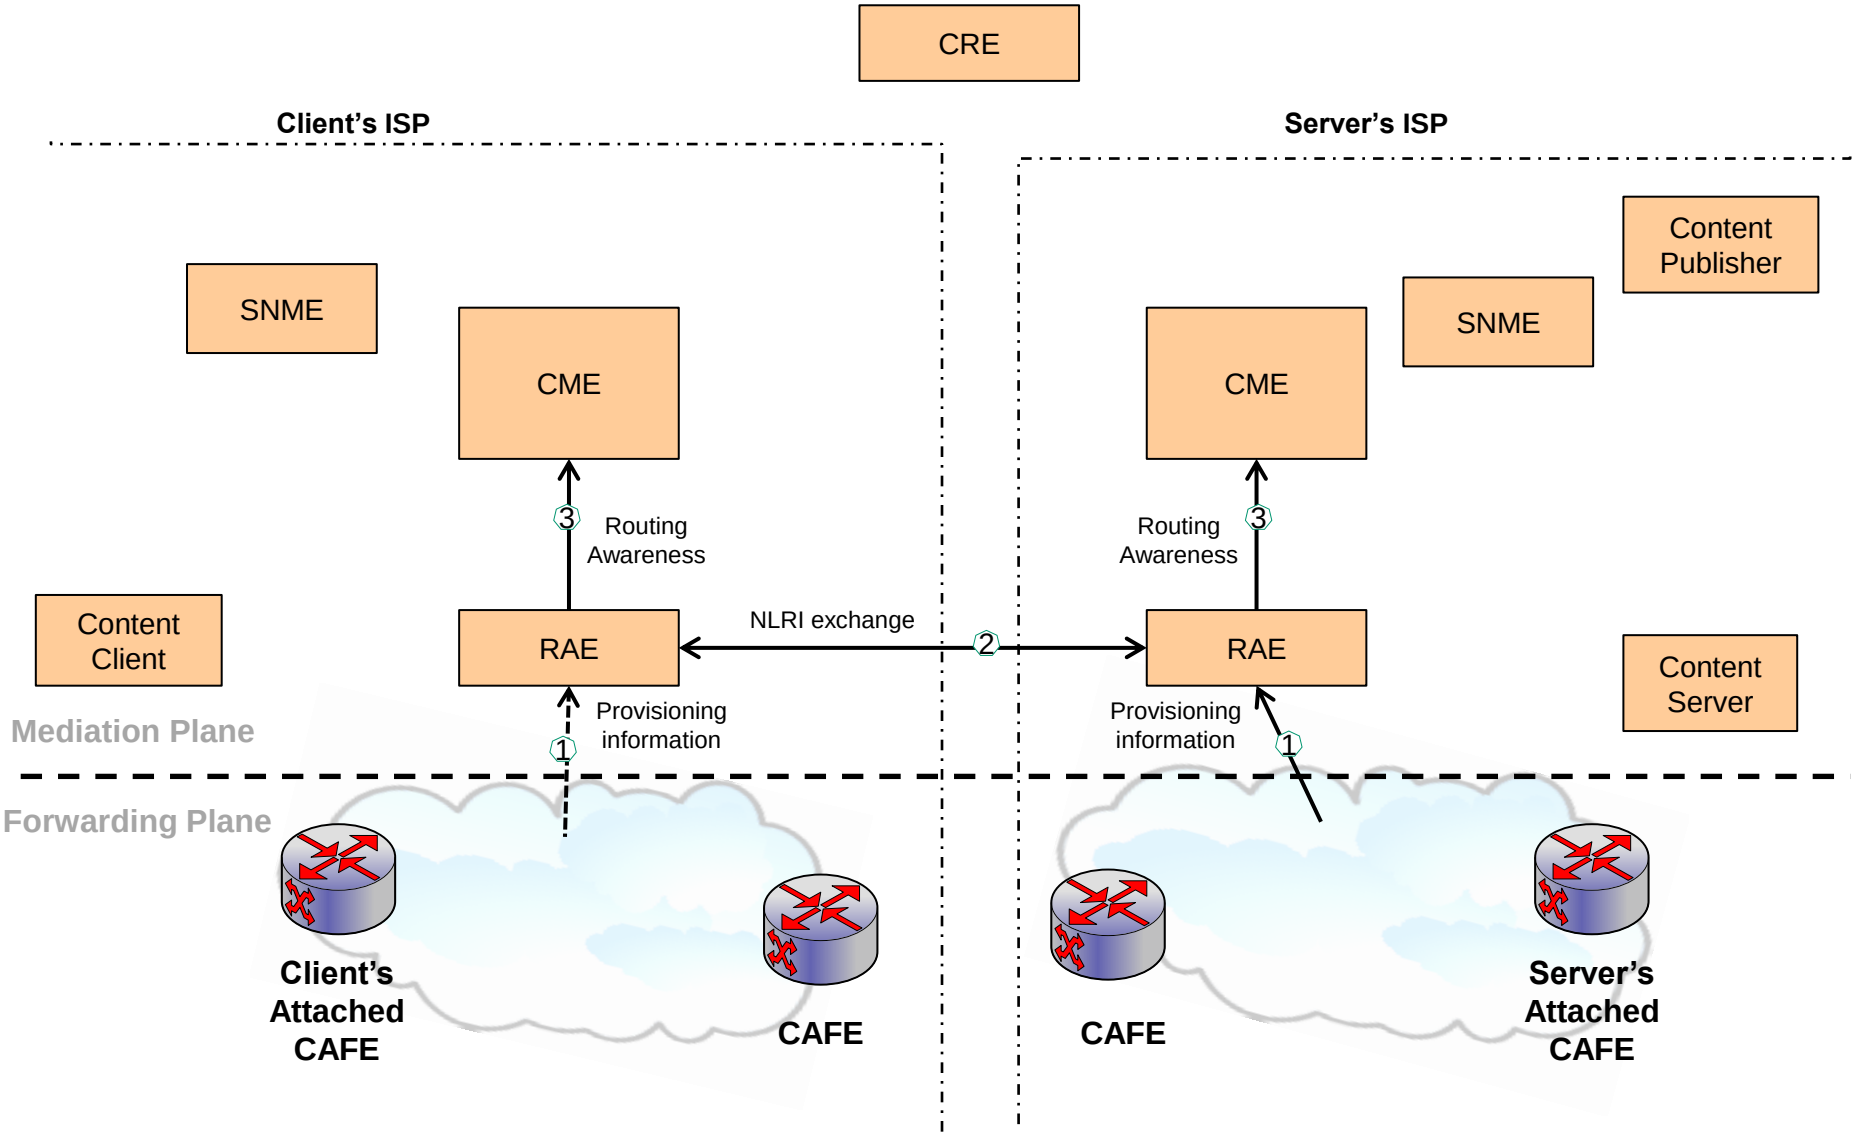

Routing Awareness (2/2)

1. RAE obtains routing information about its domain (prefixes, CoS and provisioning parameters) from provisioning tools (outside COMET's scope)

2. RAEs exchange NLRI (Network Layer Reachability Records) with peering ISPs
   - Eventually each RAE located in an ISP learns how to reach another ISP
   - Routing information is expressed in terms of prefix and list of paths
   - Path information is qualified by means of:
     - list of AS identifiers to be traversed towards given prefix (ISPs)
     - supported COMET CoS
     - QoS parameters: Max Packet Delay, Max Packet Loss and max Bandwidth,
   - QoS routing and multipath discovery
     - The RAE discovers paths with respect to requirements of COMET CoS
     - The RAE allows for more than one valid path between AS

3. RAE feeds this routing information to CME, where it is stored.

Again, this is an continual offline process, which reacts on changes in network topology or domain re-provisioning. It does not require any CME's intervention

Content Publication (1/4)

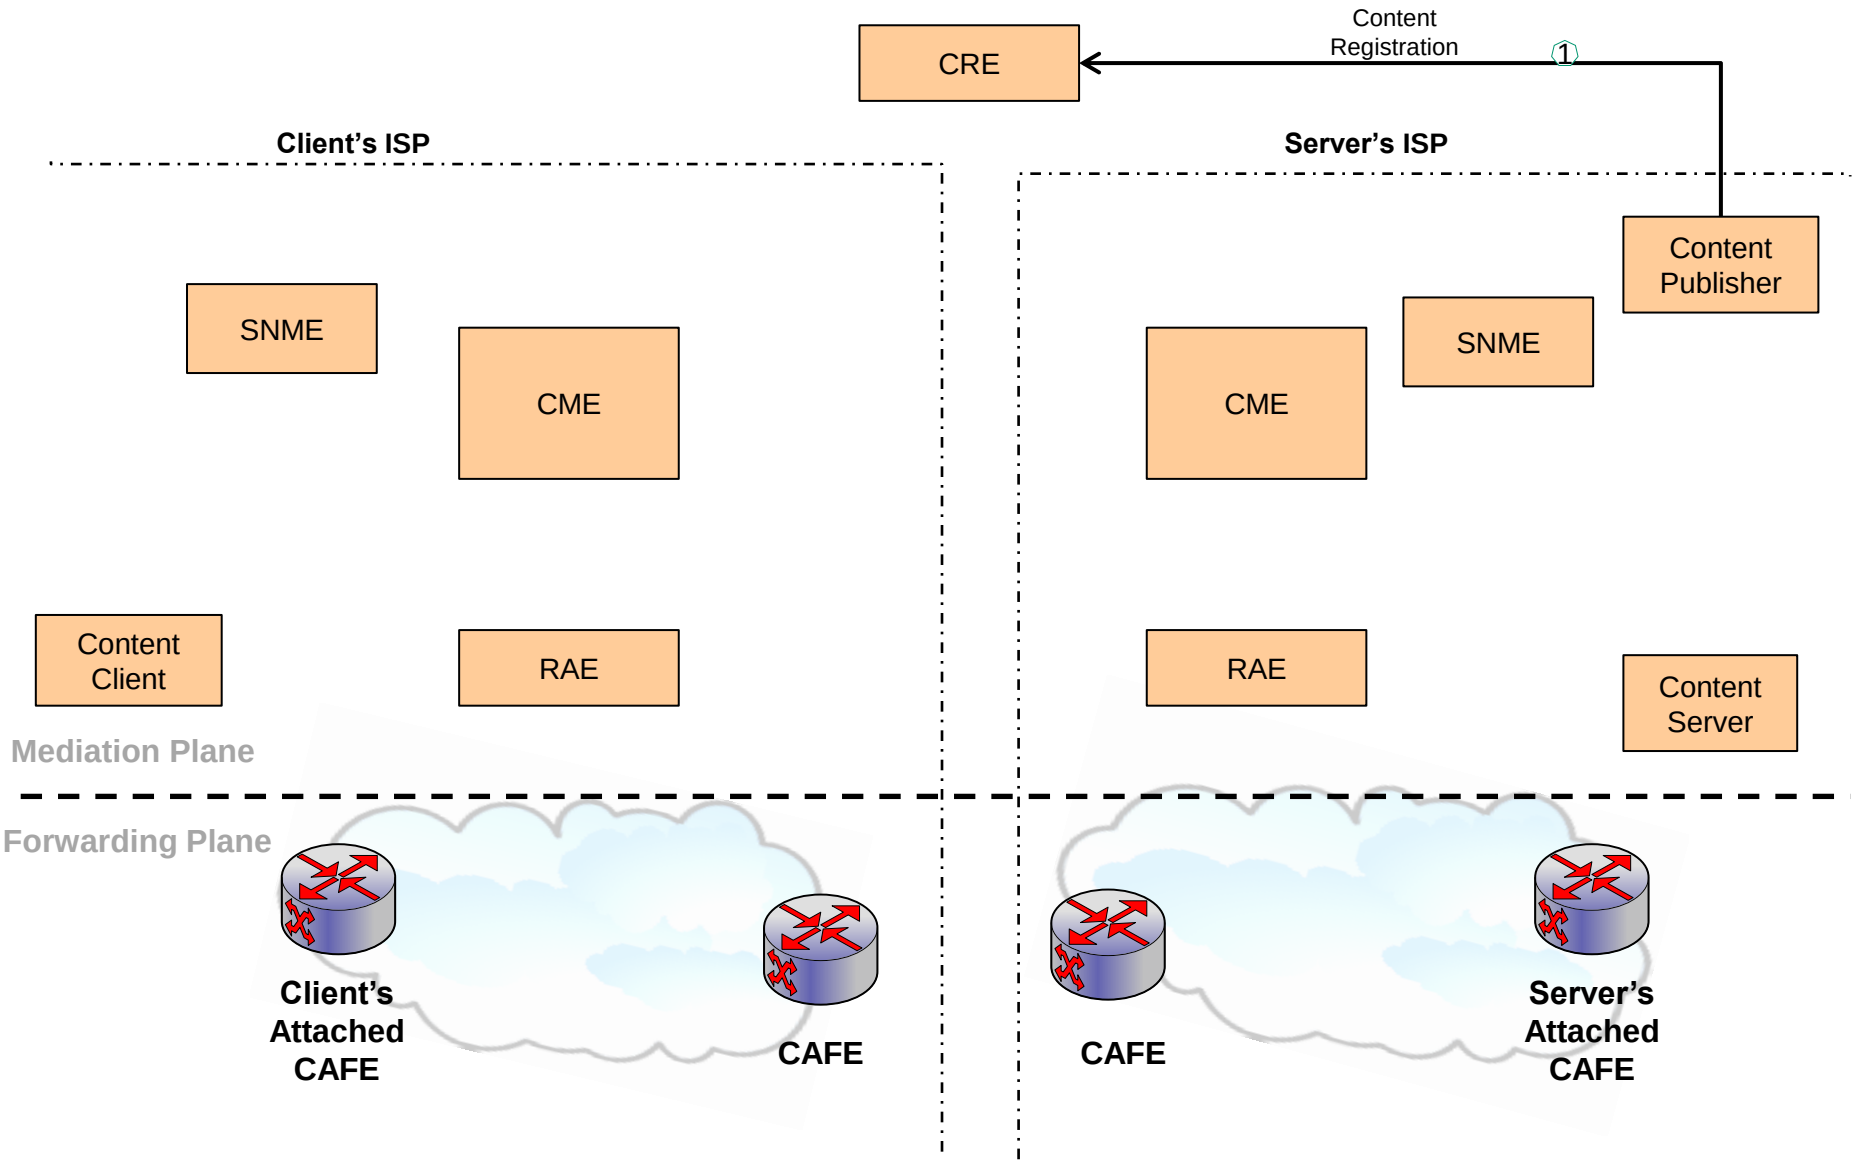

Content Publication (2/4)

1. Content Publisher sends a registration query to CRE, where it is stored for further look up

   – The CRE is the one in charge of managing the domain the content is registered to
     - i.e., to register com.gmail@John/mySong, user John will register mySong in gmail.com CRE)

   – The result of this operation is a Content Record describing the content.

   – Two Identifiers are assigned to each content
     - One Human oriented (com.gmail@John/mySong), The Content Name
     - One Machine Oriented for internal use, The Content ID

   – A Content Record contains a list of Content Sources
     - Mime Type identifying the type of content
     - Comet CoS (Premium, Better than Best Effort, Best Effort)
     - Traffic Description (Bandwidth)
     - QoS Constraints (Packet Delay, Packet Loss)
     - Application and Transport Protocol
     - Priority of the Source
     - List of actual Servers, identified by
       – IP, Path, CME ID where they are located

Content Publication (3/4)

com.gmail@john_doe/birthday_2010

CID A78F901E98AA5674A

Source 1:

MIME Type: video/mp4
CoS: Better than Best Effort
Nº of traffic descriptors: 1
Traffic descriptor: 400 kbps
Nº of QoS constraints: 1
QoS constraints: Max Packet Loss 5%
QoS constraints: Max delay 50ms
Priority 5
Application protocol: RTSP
Transport protocol and port: UDP 675
Nº of servers: 1
- Server 1:
  - IP address: 1.1.1.2
  - Path: /user/profiles/john_doe/birthday_2
  - CME ID: 1.1.1.1

Source 2:

MIME Type video/mp4
CoS: Premium
Nº of traffic descriptors: 1
Traffic descriptor: 200 kbps
Nº of QoS constraints: 2
QoS constraints: Packet Loss 10%
QoS constraints: Max delay 40 ms
Priority 10
Application protocol: HTTP
Transport protocol and port: TCP 80
Nº of servers: 8
- Server 1:
  - IP address: 1.2.3.4
  - Path: /events/2010/birthdays/john_doe/
  - CME ID: 1.2.3.4
- Server 2:
  - IP address: 5.6.7.8
  - …

Content Publication (4/4)

Important: CREs are not directly related to ISPs (*the decoupled trait*)

- An ISP does not need to own and manage a CRE
- A single CRE can manage CR for several ISPs
- A single CRE can be queried by CMEs in different ISPs.
- There will exist a hierarchy of CREs (DNS-like)
  - A CME can ask to higher levels' CREs (root) which CRE manages contents its "local" CRE is not aware of.

Content Consumption (1/6)

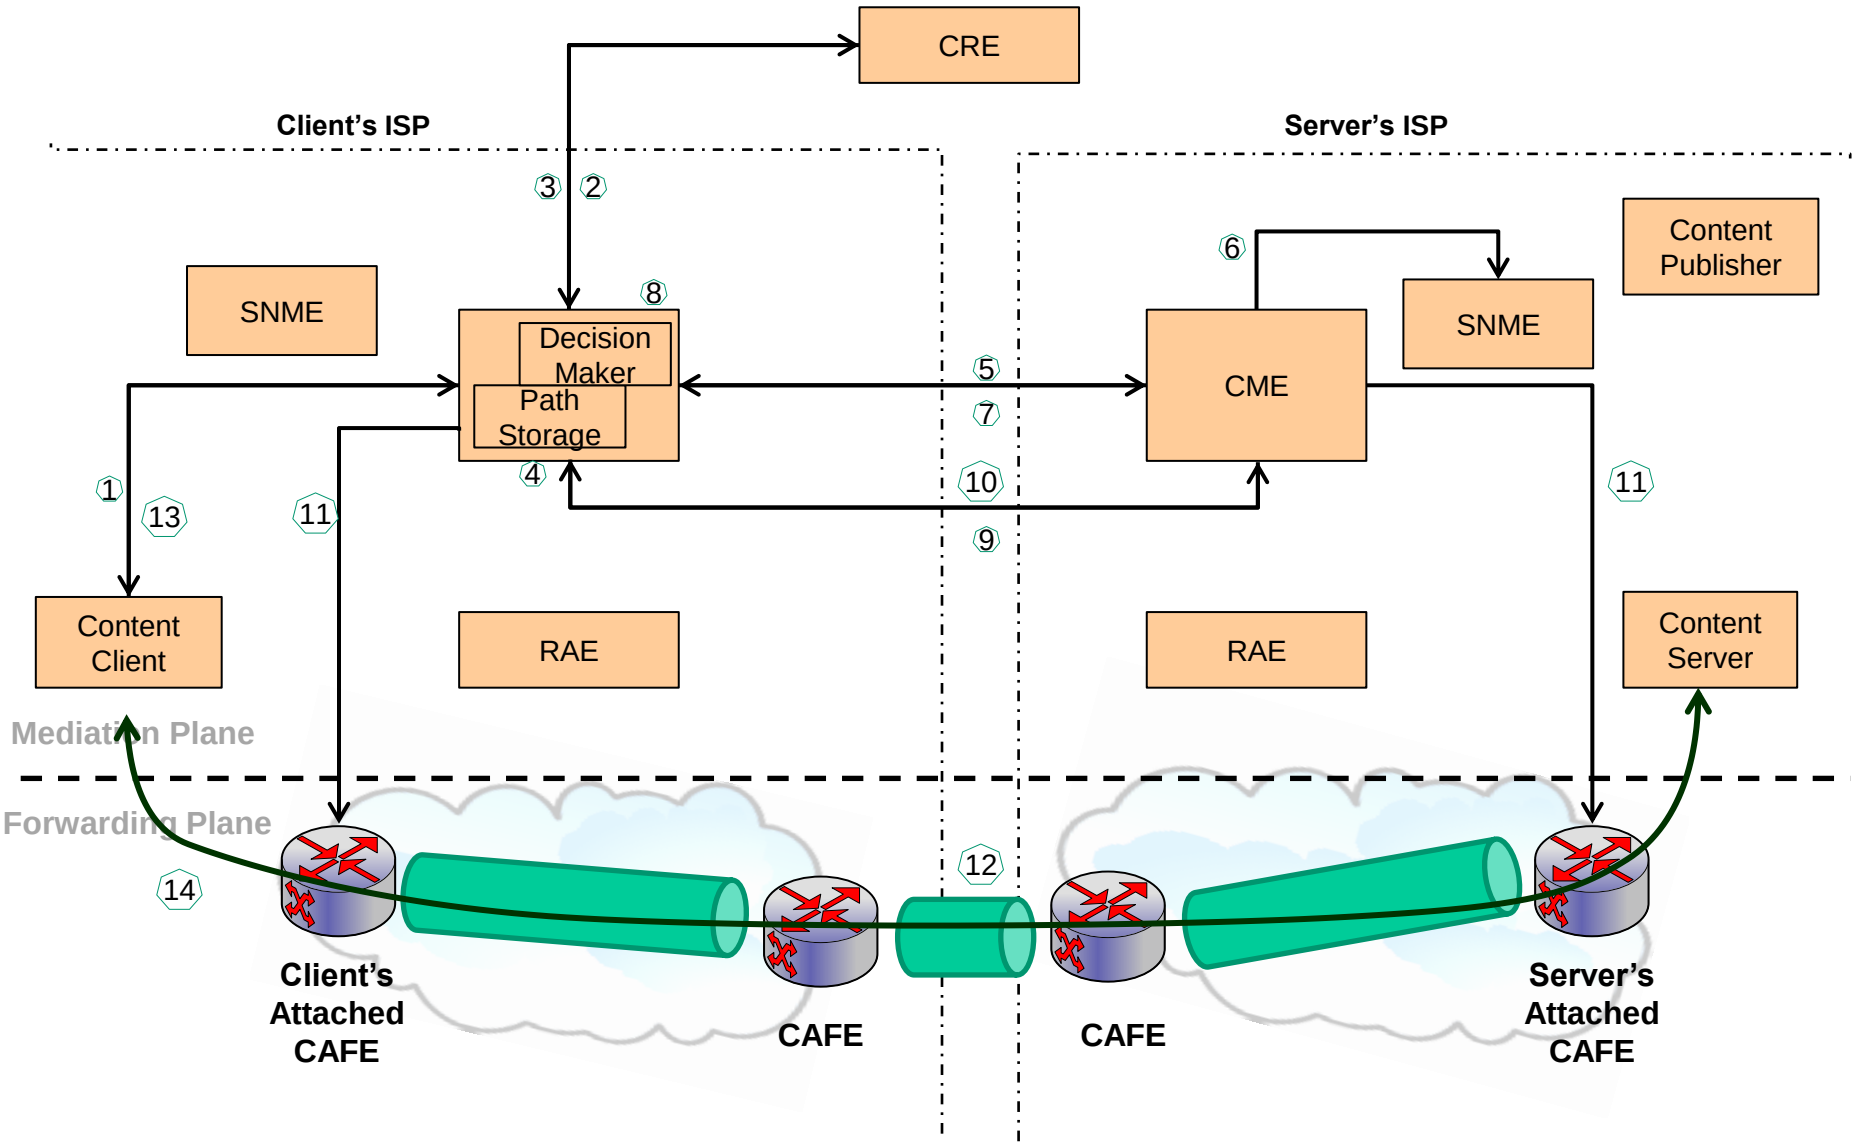

Content Consumption (2/6)

Content Request

1. CC sends a query to CME in order to consume a content (Content Name).

Name Resolution

2. CME sends a query to CRE to resolve the Content Name
   - Since CREs are “outside” the ISP, CME first has to find out which CRE can serve its query
     - Content Domain CRE, if this is already known.
     - Root CRE if CME is unaware of the CRE managing the queried domain, Root then redirects CME to the right CRE.

3. CRE looks up the queried content in its DB, retrieving the Content Record and send it back to the CME

At this moment, CME has at its disposal a list of Content Sources for the Content grouped by
- Class of Service
- Priority
- Mime Type
- Transport and Application protocols.
- Traffic and QoS constraints
- Actual Servers providing the content

Content Consumption (3/6)

Path Discovery

4. CME retrieves from its internal Path Storage the path(s) to reach the ISP of each server included in Content Record
   - For paths not stored in Path Storage, CME at client's ISP sends a query to CME at server's ISP to retrieve this information. → *Reverse Approach*

Now, CME has at its disposal

- A list of paths to each remote Server, qualified with traffic/QoS parameters
  - Each path consists of a sequence of the AS id that have to be crossed to reach the final destination,
  - Traffic/QoS Parameters are typically
    - Provisioned Bandwidth of the Path
    - Packets Loss
    - Packet Delay
- The List of available Content Sources, as described above

Content Consumption (4/6)

Server Awareness

5. For each server in the Content Record, Client's CME queries remote CME at servers' ISP, in order to retrieve their load status

6. Each remote CME searches for the server status in local SNME.

7. Each remote CME forwards the status information of its servers to CME of client's ISP.

At this point in time, Mediation Controller at the client's CME has gathered information about Content Sources, Qualified Paths to them and Servers' Status

Decision Process

8. The consolidated information is assessed by Decision Maker in the CME, which (hopefully) returns the best path option.
   - Path is still described as a list of AS Identifiers.
   - It has to be translated into a list of CAFE's identifiers
     - CAFEs in each ISP configured to forward packets from that ISP to the next in the list.

Content Consumption (5/6)

Path Configuration

9. Local CME at Client's ISP sent the path info to the remote CME at Server's CME
   - Path in AS form + Client's attached CAFE.
   - Remote CME adds the Server's attached CAFÉ
   - Finds the CAFE linking the remote ISP with the penultimate ISP in the list
   - Queries that ISP to resolve the list of remaining ISP after having deleting itself
   - An so on recursively

10. The translated Path is returned to the CME at Client's ISP

11. Each CME in both client's and server's ISP configures CAFEs involved in the path.
   - Only client's attached CAFE and server's
   - With the list of forwarding CAFEs to be traversed to reach the other end.

Content Consumption (6/6)

12. Finally, a tunnel is ready for content retrieval
    - Binding client's attached CAFE to server's

13. CME of client's ISP sends content connection parameters to CC
    - Mime Type
    - Application Protocol
    - Transport Protocol
    - Server IP and Port

14. The client downloads the content through the tunnel

    - When a packet is received at client's attached CAFE or at server's the list of CAFEs to be traversed is added to the packet
      - Subsequent CAFEs in the path know how to route the packet by looking the CAFEs' list up

COntent Mediator architecture for content-aware nETworks