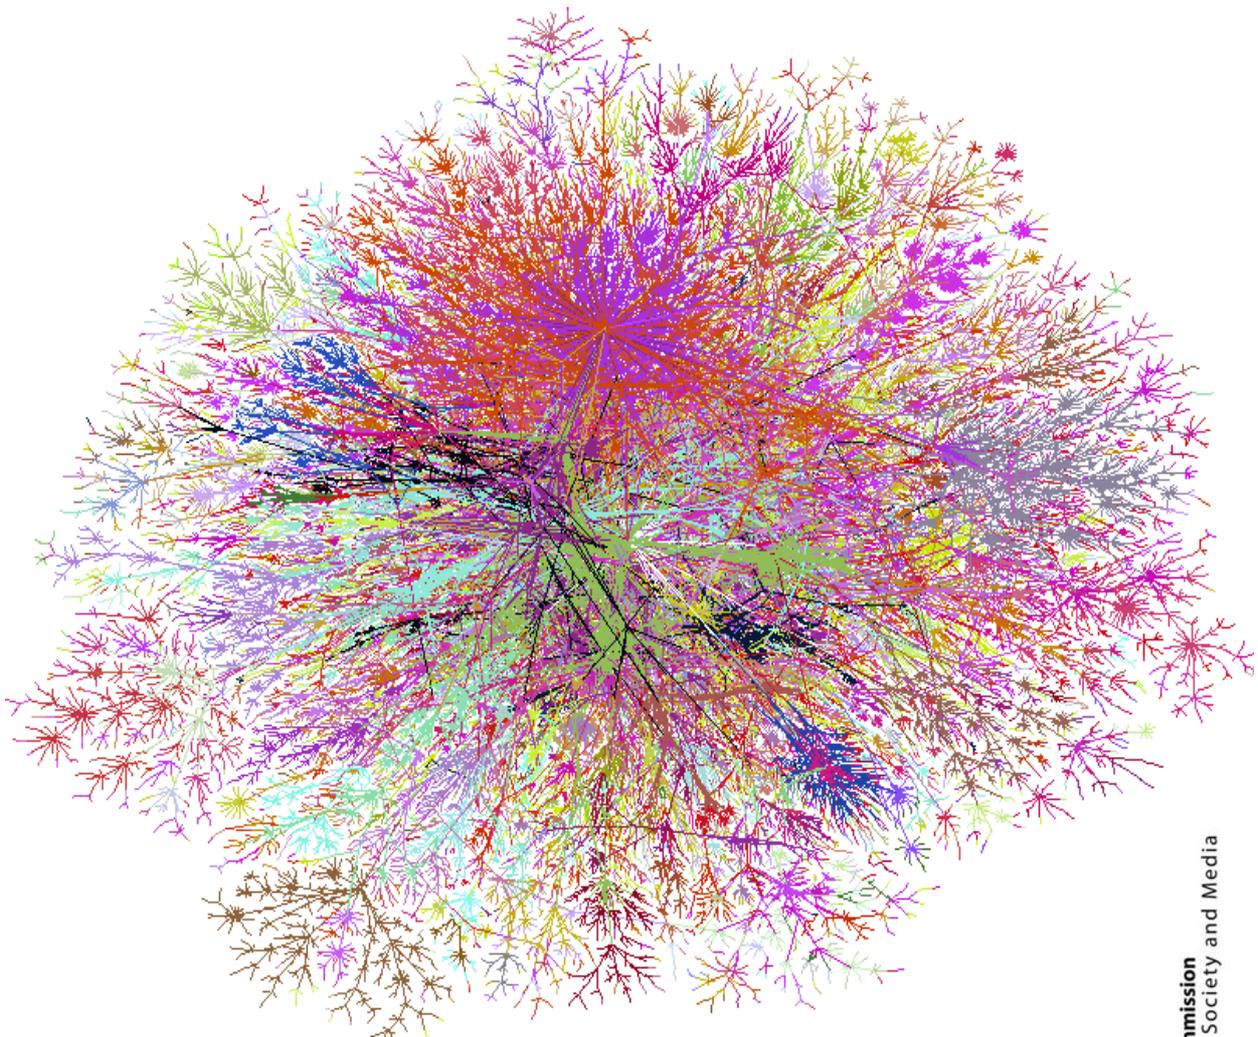




Why do we need a Content-Centric Future Internet?

Proposals towards Content-Centric Internet Architectures



Created by the Future Content Networks Group
Prague, May 2009



European Commission
Information Society and Media



Why do we need a Content-Centric Internet?

Proposals towards Content-Centric Internet Architectures

Executive Summary

The aim of this document is twofold: firstly, to report and analyse the main reasons, which support our claim that the Future Internet will be “Content-Centric” and secondly to define two alternative solutions for a Future Content-Centric Internet Architecture following an evolutionary and a clean-slate approach.

The starting point of our discussion is the reasonable hypothesis that Future Internet will mainly simplify the usability, increase the efficiency, secure the privacy and enhance the media experience of the users (enhanced mobility, really broadband & flexible communications, immersion, enhanced interaction, involvement of all senses and emotions, navigation). New ways of media creation and consumption will emerge, aiming to cover the different human needs and preserve the revenue generation of the various stakeholders. Moreover, new content types will appear, which together with efficient handling, delivery and protection of the content (i.e. static or dynamic, pre-recorded, cached or live) will be the Future Internet cornerstones. Thus, the content/media and its efficient handling are (in) the heart of the Future Internet.

Taking into account the fact that the current Internet cannot efficiently serve the increasing needs and the foreseen requirements, two Content-Centric Internet Architectures are proposed: a “Logical Content-Centric Architecture”, which consists of different virtual hierarchies of nodes with different functionality and an “Autonomic Content-Centric-Internet Architecture”, which relies on the completely novel concept of the “content object”.

Yet, the major objective of this position paper is to initiate a debate between all the interested stakeholders with respect to the following three fundamental arguments:

1. Will the Future Internet be Content-Centric?
2. How a potential Content-Centric Internet Architecture would look like?
3. Which design principles and requirements would govern such Architecture?



Table of Contents

1	Will Future Internet be Content-Centric?	4
1.1	What will Future Internet bring to humans?	4
1.2	Why should Future Internet be content-centric?	7
2	Current Internet Operation/Limitations	9
3	Future Internet Design Principles and Design Requirements	11
3.1	Future Internet Design Principles	11
3.2	Future Content Centric Internet Design Requirements	12
4	Future Content-Centric Internet Architecture	14
4.1	Logical Content Centric Architecture	14
4.2	Autonomous Content Centric Architecture	17
5	Conclusions	19
6	Acknowledgements	19
7	Annex - Future Network For a examples	20
7.1	Digital Living Network Alliance (DLNA)	20
7.2	IP Multimedia Subsystem (IMS)	20
7.3	Open IPTV Forum	21
7.4	Conclusion	21
8	List of Contributors	22
8.1	Editorial Team	22
8.2	List of Contributors	22

Notes:

1. The figure at the cover page is the "Map of Internet: colored by IP addresses" by W. R. Cheswick
2. The FCN logo is the combination of the Future Content Networks Acronym, surrounded by one of the four water clusters, singled out by F. H. Stillinger, "Water revisited," Science (1980) 451-457. It is based on the metaphor that FCN, just like the water surrounds the world.

1 Will Future Internet be Content-Centric?

The Internet is clearly evolving to a Media Web, where both creation and consumption of audio-visual content is expected to continue to rise in the foreseeable future, together with the quality of that content (resolution, frame rate, color depth, stereoscopic information). Besides these clearly observable trends, we can expect changes in the way the content is published that will open up new possibilities for content creation and consumption. These changes will create new challenges for the network infrastructure and the services offering.

Yet, the first fundamental questions that we try to analyse in this document are related with the orientation of the Future Internet itself. They may be expressed as:

- ***What will Future Internet bring to humans?***
- ***Why should Future Internet be content-centric?***

These fundamental questions are not only difficult to be answered but also different answers may lead to different designs and orientations. Some of them may not be excluded, leading to parallel viable Internets, while others may not be visionary enough, leading to artificial Future Internet Architectures, which may be better than today's Internet, but not "the one and only". In the next sections we try to give a first reasoning for our selection to support a Future Content – Centric Internet (FCCI).

1.1 What will Future Internet bring to humans?

In answering this question, it is important to consider the nature of digital content today and in the near and longer future. Moreover, we should consider the Future Internet from both a socioeconomic and technological point of view.

Starting from the socioeconomic point of view, we expect that the Future Internet will bring new ways of multimedia content creation and consumption, aiming to cover the different human needs and preserve the revenue generation of the various stakeholders. Over the last few years, we have seen huge increases in the popularity of professional content delivered using Internet networks. However, professional digital content is a relatively low-margin product, because the cost of acquisition/creation can be as high as 40% of revenue. Multimedia broadband communications and user-generated content (either pre-recorded or live) may increase rapidly the network traffic and be among the major revenue sources for the network infrastructure providers. Moreover, service providers will take advantage of the new features and opportunities offered by the Future Internet in order to provide new added-valued services. On the other hand, human factors, such as the increasingly distributed nature of our families and greater demands on personal time, also support the argument that people will rely on the Future Internet to meet more of their higher order social needs.

In short, if we try to go down to the very basics, today's Internet bring content. It may be text, voice, audio-visual or multimedia; it may be pre-recorded or live; it may be professional or user generated. Content (or stream session) finding and retrieval are the first steps of utilising the content. Information is the result of content mining, combination and aggregation. Services are the result of content or information manipulation and utilisation. Security and privacy are applied (not only, but in a large degree) on the content. At the very end, today

people use the Internet due to its content (stored or live) or the information and services that they may get out of that.

In an evolution way, towards the question “*What will Future Internet bring*” a short answer may be: “*Future Internet will bring new experiences and communications features*”. The Future Internet should accommodate services, methods, procedures and techniques that would enable the creation and consumption of new experiences (either professional or user generated) and seamless, personalised, multimedia communications, across a grid of network constituents, such as personal area networks, body area networks, home networks, fixed access networks, mobile access networks, metro networks and core networks.

This answer includes new services, new communication means, new applications, new streaming capabilities. Yet, if we try to go again down to the very basics, all these end-up at least to *New Media*. In the Future Internet, in order to realise the “*new experience*”, we will need new media and content that may capture more senses (visual, audio, touch, emotions, even smell when recreation problems are solved). Figure 1 demonstrates these new media-based paradigms.



Figure 1: Immersive Future Media paradigm¹

Moreover, emerging technology is expected to make it possible for constituting components of audio-visual content to survive the digitization and publication process, or at least be resurrected again:

- **Automated audio-visual analysis techniques** can segment audio-visual content and resurrect the composing components. The technology is currently used to optimize media encoding or to allow media search, and will required in the short to mid-term to decompose into constituting components the audiovisual content that is the result of past and current capturing technology.
- In the long term, we will depend on **component-preserving capturing technology**: advanced capturing systems using stereoscopic cameras, camera arrays, time-of-flight cameras or 3D scanners preserve three-dimensional information and thereby the components and their spatio-temporal relationships.

¹ Figure provided by P. Daras, CERTH/ITI

- When 3D modeling is used to create audio-visual content, ***all the information about constituting components and the composition process will be readily available***. When photorealistic 3D content creation technology that produces content that is indistinguishable from captured content comes within reach of prosumers, we can expect a sharp increase in output of this type of content.

The availability of the constituting components and their spatial and temporal relationships, rather than an opaque stream of pixels and audio samples, opens up new opportunities for content creation and consumption:

- Re-use of components from existing content for the creation of new audio-visual content becomes much less cumbersome, allowing fast and easy media mash-ups
- On-line collaboration to edit audio-visual content prepares the way for a true media Wikipedia
- Personalisation enters a new stage, evolving from a selection of prepared content to a just-in-time composition
- The insertion of stored audio-visual content into real-time communication is facilitated.
- The combination of captured audio-visual content with synthetic 3D content creates exciting mixed-reality experiences

The Future Content Centric Internet will offer the end-user enhanced immersive real-time experiences that can address both content consumption scenarios and communication & collaboration scenarios. The realisation of enhanced immersive real-time experiences involves the ad-hoc, on the fly generation/reconstruction of semantically enriched 3D augmented/virtual worlds by the integration of professional and user-generated media components that are collected from various sources in and around the network (Figure 2). Taking into account the spatial and temporal aspects during the just-in-time integration creates an orchestrated immersive media experience.

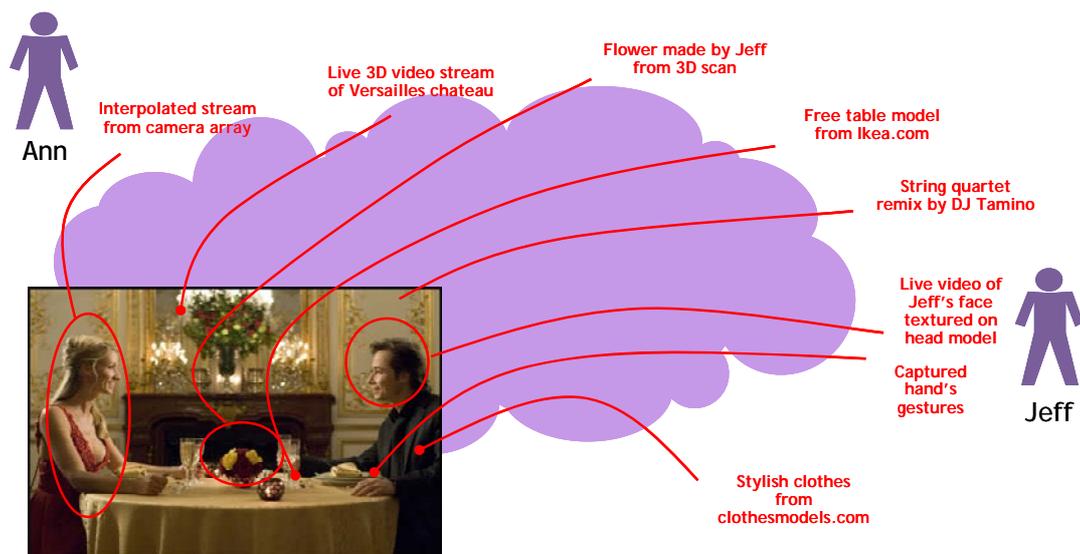


Figure 2: On the fly generation/reconstruction of semantically enriched worlds²

² Figure provided by Jan Bouwen, Alcatel-Lucent

1.2 Why should Future Internet be content-centric?

Of course, Future Internet will not only bring new media and content. Attention should be also paid to the presentation, communications, services, infrastructure and privacy, which are essential building blocks of the FI. Yet, the content (new or traditional, broadband or narrowband, professional or user-generated, pre-recorded, live or interactive, broadcasted or personalised) is what the users will receive/enjoy from the Future Internet. Starting from the reasonable assumption that the content and the content representation (current forms and new media) is the basis of the Future Internet, and following a bottom up approach, we may assume that if the Future Internet is content-centric, it can efficiently handle the content and this will lead to a new generation of the Internet.

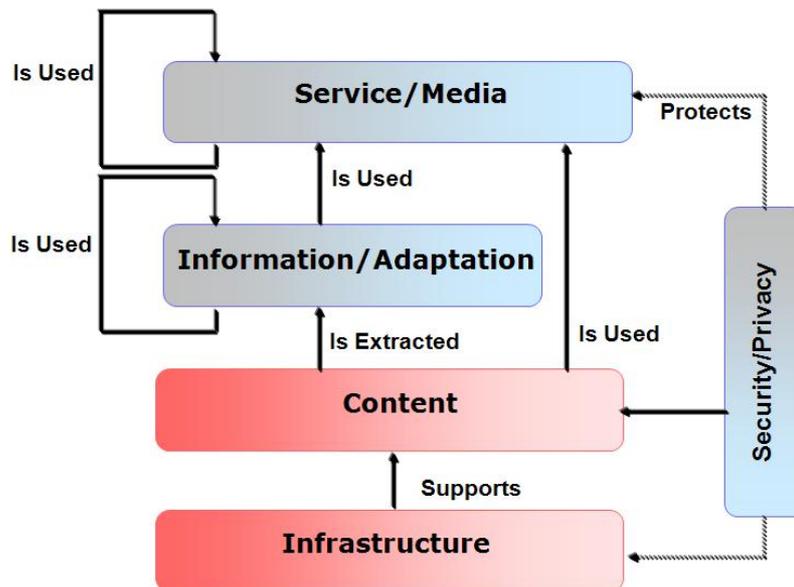


Figure 3: Future Content-Centric Internet components interrelation

In more details, we may define (Figure 3) as:

- **Content** is any type and volume of media. Content may be pre-recorded, cached or live, static or dynamic, monolithic or modular. Content may be combined, mixed or aggregated to generate new content and **Media**. It may vary from a few bits (e.g. the temperature that a sensor has measured) to interactive multi-media sessions and immersive complex and multi-dimensional virtual/real worlds' representations.
- **Information** is the product of a number of functions applied to the *content* or recursively to the *information*. By combining, mining, aggregating *content* and pieces of *information*, new *information* may be extracted or generated.
- **Service** is the result of a set of functions applied to the *content*, to pieces of *information* or recursively to *services*. By (manually or automatically) handling, managing, combining, personalising, adapting *content*, *information* or *services*, new *services* may be composed or generated.
- **Infrastructure** (both private and public) will consist of transport, storage and processing functions in a distributed manner. This cloud offers the opportunity to deal with active



content objects, rather than unstructured bitstreams in the network. Thus, intelligent content-centric infrastructure may efficiently support content and its derivatives (information and services).

- **Security and Privacy** will be a property of *content, information, services* and Infrastructure, allowing much more efficient control over *content* objects.

Even though these considerations may look oversimplified and debatable, one may admit that the better the content manipulation is, the better the Future Internet can be. This is why we define the Future Internet as Content-Centric.

If Future Internet is content-centric, not only will it enable the next generation of personalised services, delivering pre-recorded or live content and information, but it will also open up a new world of high quality audiovisual communications, which requires content to be captured, streamed and presented consistently and with minimal delay.

2 Current Internet Operation/Limitations

Today, the vast majority of Internet usage concerns content and services discovery & retrieval, content delivery and streaming and Web services access, where the user cares only about the content or service itself and is oblivious to their location. That is, the user knows that s/he wants news from the BBC, videos from YouTube or weather information, delivered in suitable quality and format, but s/he does not know or care on which machine the desired data or service resides, as soon as security and privacy are guaranteed.

The above functionality is realised by the current Internet Architecture as shown in Figure 4. It consists of the following types of nodes:

- Content Servers or Caches** (either professional or user generated content and services),
- Centralised, decentralised or clustered **Servers**, including *Search Engines*, but also *Supporting Servers* e.g. DNS servers, AAA servers, DRM servers, etc.
- Core and edge **Routers** and **Residential Gateways** (represented as R1 to R5) and
- Users connected via fixed, wireless or mobile **terminals**.

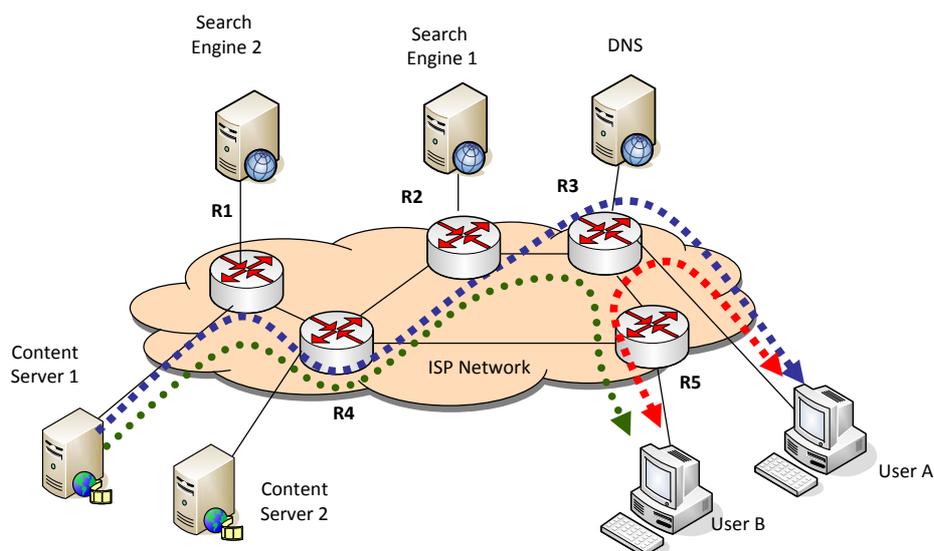


Figure 4: Today's Network Architecture, Content Discovery, Retrieval and Streaming

The initial step is **Content Discovery by the Search Engines**: the Search Engines crawl the Internet to find, classify and index content or services. Alternatively, users may publish content and manually inform the search engine. The second step is **Content Discovery by the User**: if the user does not know where the content resides, s/he queries a Search Engine and gets as feedback a number of URLs, where the content is stored. The last step is **Content Delivery/Streaming**: the user selects a URL and the content is delivered or streamed to her/him. Alternatively, in case of live communications services (e.g. VoIP or video conference), User A and User B are communicating using their IP addresses as reference.

In the scenario shown in Figure 4, if both User A (UA) and User B (UB) ask for the same content to the same Search Engine, they will both get as an answer that the content is stored at Content Server 1 (CS1). Then, assuming no caches or mirrors, the content will be delivered via the routers' (Rx) path: e.g. CS1-R1-R4-R2-R3-UA and CS1-R1-R4-R2-R3-R5-UB, respectively.

The above schema works for current applications and usage, and will continue to do so provided there is sufficient resource in the system to deliver it. “Resource” may mean bandwidth capacity in a given link or it may mean the capacity to route a packet of data with a sufficiently low delay. But what happens as:

- billions of devices become connected; capable of being interrogated and providing information that can be aggregated into services?
- users demand image resolutions in video that require bandwidths greater than can be supported over the typical link lengths using the current physical link layers
- more and more users conduct delay critical real time video and audio communications using the Internet

These changes will only be supported in the current Internet through massive investment, and even then the architecture may exhibit unstable characteristics. An intelligent evolution of the Internet architecture will lead to much more efficient use of the available resource (bandwidth, routing capacity) and provide a business environment that encourages investment.

Consider then, some changes to this schema that would make better use of the available resource. Significant improvements can be thought in order to efficiently deal with Future Media (Figure 4). More specifically, what if:

- the **content could be stored/cached closer to the end users**, not only at the end-points as local proxies, but transparently in the network (routers, servers, nodes, data centres) then content delivery would have been much more efficient.
- the **routers could identify/analyse what content is flowing through them** and are able to replicate it efficiently, the search engines would gain much better knowledge of (even the streaming) content location and provide information even on “live” video streams. For example, if R2 could report to Search Engine 1 that “channel 5” content flows through it, when User B asks for “channel 5” Search Engine 1 could reply R1 instead of CS1.
- the **network could dynamically identify what is the best path to the user** (less congestion, lower delay, more bandwidth), it would have provided a better way to deliver the data. For example the path R2-R4-R5-UB may be much better than R2-R3-R5-UB. This also includes innovative P2P streaming and dynamic constructed overlays.
- the **content could be interactively adapted**, not only statically based on the network and terminal capabilities, but also based on the interactive content selection by the user (e.g. instantly changing the point of view, zoom-in/zoom-out at a streaming session), the user experience would be much better.
- the **content could be selected and adapted to the context**, the user would have a much easier life in the content-centric Internet. Based on entering a living room, a phone TV session could transfer to the big screen and adapt to the resolution offered there. This likely involves using a completely different access and transport network as well.
- the content **could be active instead of static**, it would be much easier to be able to find out about best presentation, user rights, network path etc. It would mean that besides media encodings content objects would contain methods which implementations would allow them to control their rendering as well as their transport.

3 Future Internet Design Principles and Design Requirements

Future Internet may be realised in a vast number of alternative ways³. Yet, in this section we try to define the main principles and requirements that will lead the design of the Future Internet especially from a content perspective. The section is organised in two parts, one general pointing to design principles shared with content-centric networking requirements and one specific.

3.1 Future Internet Design Principles

From a technological, content-centric viewpoint, four main principles emerge towards the design of the Future Internet:

- **Support multiple and new Business Models.** The FI should be capable of supporting flexible business models. Multiple stakeholders (individuals, communities, small businesses, large enterprises) should be able to participate in an open environment that supports and encourages innovation and participation without barriers. This principle includes the possibility to build a Future Internet that could host various embedded architectures, one of them content-centric.
- **Simplicity:** In current Internet the complexity belongs at the edges, and the IP layer of the Internet remains as simple as possible. Complex systems are generally less reliable and flexible. On the other hand, complex problems sometimes require complex solutions and the FI will be providing non-trivial functionality in many respects. Thus, when designing the Future Internet, the famous quote by Albert Einstein may be adopted: "Make everything as simple as possible, but not simpler". It will anyway be a continued requirement to make the usage of network functionality simple and robust.
- **Sustainability, Scalability and Robustness.** The FI must be designed as a sustainable network, offering built-in support for energy efficient solutions, being flexible enough to continuously evolve, and to develop and extend in response to changing societal requirements. FI should be able to serve a very large number of entities (scalability), maintaining its usable operation ratio (availability) and can easily recover if faults occur (reliability). It should contain and probably expand the principles of any management system such as Fault, Configuration, Accounting, Performance and Security (FCAPS)⁴.
- **Loose Coupling.** As things get larger they often exhibit increased interdependence between components. The FI should be built allowing for loose coupling, since loosely coupled systems have more flexibility in time constraints, sequencing, and environmental assumptions than tightly coupled systems.

³ R. Bush, D. Meyer, "Some Internet Architectural Guidelines and Philosophy," RFC 3439, Dec. 2002

⁴ W. Willinger, J. Doyle, "Robustness and the Internet: Design and evolution", 2002, http://www.maoz.com/~dmm/complexity_and_the_internet/robustness_and_the_internet_design_and_evolution.pdf

3.2 Future Content Centric Internet Design Requirements

Outcomes of an initial analysis on the FCCI design requirements for full media experience led to the main requirements outlined below⁵. These requirements are drawn from a media-centric analysis and consequently they shift the emphasis to content and tussle driven design, findability, trust and flexibility.

- **Content-Centric Engineering:** Assuming that the Future Internet Architecture will be “content-centric”, then efficient content engineering should be at the centre of the FI design requirements. Basically, the design of FCCI should include new ways of storing, coding, enriching, finding and rendering content, where the content is treated as having meaningful semantic connotations rather than simply a set of pixels that have been encapsulated in packets for transmission over the network. There are several design requirements related to the content engineering aspects of FCCI.
- **Name Resolution and “Findability”.** Users are typically interested in accessing a particular piece of content, rather than where the content is stored or the name, URL or IP address of the server that hosts it. Therefore, the focus should be on the content and the information that may be extracted by content-mining/pre-processing; neither on the content location nor on the network topology. **“Findability”** contains two operational primitives: Find and Register. Users should be able to request named data from the network by using the “Find” primitive, while content providers can publish a data object, which will be served to the users by using the “Register” primitive. A FCCI should allow for ease of content advertisement, discovery and retrieval.
- **Content-Centric Routing.** The FCCI should be able to support content-centric routing, which allows users to fetch particular pieces of information from the most convenient location or locations based on the user’s requirements (e.g., minimizing download time or minimizing latency for real-time content). In any case, users should be able to deal with information without concerning themselves with the location of that information. Such an approach would lead to multiparty-to-multiparty information dissemination rather than traditional point-to-point conversations.
- **Flexible Content Business Models.** The FCCI should be capable of supporting flexible business models where multiple stakeholders can participate in an open environment that supports and encourages innovation and participation by users (individuals or communities), ISPs, content and service providers, public and private organizations and regulators. This design requirement is possibly the main differentiating point between NGNs and the envisioned FCCI^{Error! Bookmark not defined.}. FCCI should also be designed for tussle. The term “tussle” describes the clash of interests between Internet stakeholders⁶. For example: ISPs must interconnect to provide global connectivity, but they may be competitors of one another; users may wish to make use of services from any provider, but ISPs may want to restrict their customers. Some key principles to be ensured include: freedom of choice, openness of access, removal of barriers to innovation and means of establishing trust.

⁵ Future Media and 3D Internet Task Force, “Future Internet and NGN: Design requirements and principles for a Future Media and 3D Internet,” Networked Media Unit, February 2009

⁶ D. Clark, J. Wroclawski, K. Sollins, R. and Braden “Tussle in Cyberspace: Defining Tomorrow’s Internet” In Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols For Computer Communications, SIGCOMM ’02, New York, NY, 347-356.

- **Trustworthiness of content and media services.** One of the FCCI challenges is to provide services and their content securely between all users, guaranteeing the privacy of each participant in a media transaction and securing networks against breakdown and malicious attack. Further user confidence criteria must be satisfied: e.g. trusted services and applications, full user control and awareness of what data can be available, rights-management/protection across the entire value chain for complex, effective controls (e.g. parental controls) over access to unsuitable content; trust measurement and management methods, e.g., reputation.
- **Choice of source and presentation.** The FCCI should provide choice as to how and from whom a user fetches a particular piece of content or information. For content or information stored at multiple providers, the FCCI should allow choosing from which the information is fetched. In some cases, where to fetch a particular piece of information may be a trade-off between interests of multiple players. Content should be available in various presentation formats, The appropriate format should be adapted to the user's preferences and situation. The FCCI should also support the automatic allocation of resources (e.g., network capacity) appropriate for a particular content delivery and presentation.
- **Decentralised Self-organization.** The proposed content-centric network architecture should not require large-scale manual management or configuration (e.g. of routing information, etc.), or even global coordination. In this way it will not be subject to the vagaries of large-scale attacks or mis-configuration. On the contrary, the primary focus should be to provide for a decentralised, secure self-organization of named content.

4 Future Content-Centric Internet Architecture

In this section, we describe two approaches towards FCCI architecture. The first sub-section offers a logical content-centric Architecture that spans from the infrastructure to the applications layers, following a more evolutionary than a clean-slate approach. The second sub-section concentrates on the content aware functionality and proposes a visionary forward-looking “object” approach, which may be the Future Internet content, information and services organisation.

It is important however to understand that any architectural evolution will not be successful if it is proposed in isolation. Today, several frameworks of industrial recommendations and standards exist, which have been established to promote agreement and interoperability between many of the organisations, which will ultimately be the major commercial players of the Future Internet. In Annex, we describe some of the most relevant frameworks and the extent to which they are already addressing content-centric architectures.

4.1 Logical Content Centric Architecture

In the Future Internet, intelligence will be moved in the network and transform the Network Architecture of Figure 4, into a content-centric architecture as shown in Figure 5. Assuming a progressive, rather than aggressive evolution towards Future Internet, the architecture may consist of different virtual hierarchies of nodes (overlays, clouds or virtual groups of nodes), with different functionality. In the figure only 3 layers are shown; yet, this model may be easily scaled to multiple levels of hierarchy (even mesh instantiations, where nodes may belong in more than one layers) and multiple variations, based on the available level of information and service delivery requirements and constrains.

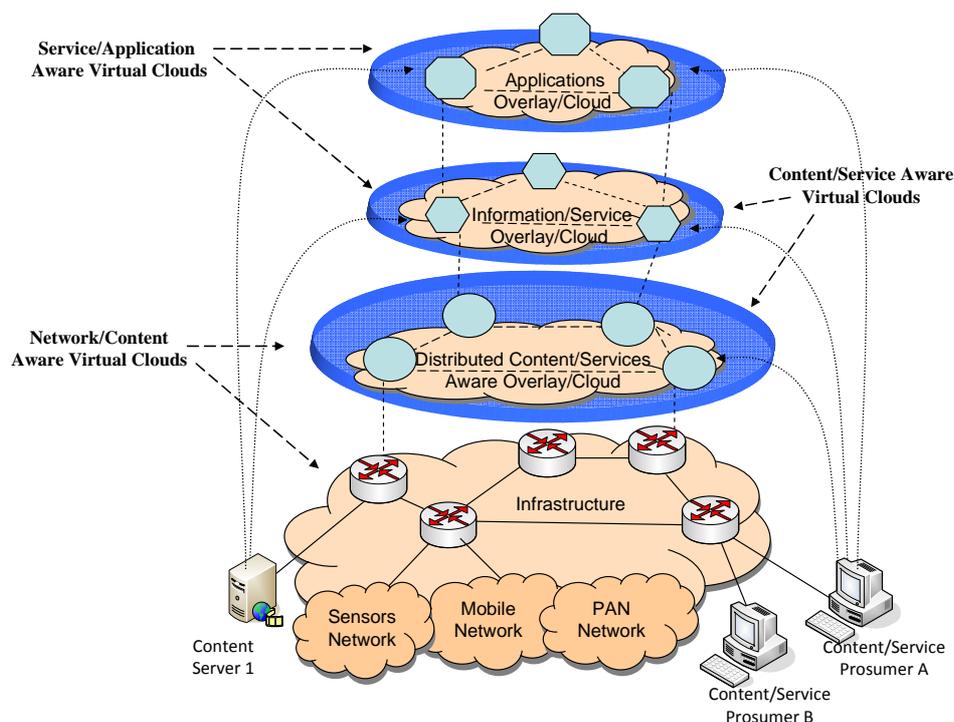


Figure 5: Logical Future Content-Centric Internet Architecture

At the lower layer, it is the **Service/Network Provider Infrastructure**. Today's existing legacy network nodes (core routers, switches, access points, residential gateways) will be the majority of the network nodes for a number of years. This Service/Network Infrastructure is the service, ISP and network provider network infrastructure, which will consist of nodes with limited functionality and intelligence. Users are connected to the infrastructure as content producers and consumers (Prosumers). Content will be routed, assuming basic quality requirements and if possible cached to some degree in this layer. Progressively this overlay will be reduced or even eliminated.

The Infrastructure should on the one hand, hide all unnecessary complexity (e.g. physical network topology, mobile terminal handover) and on the other hand provide all the necessary information, so that more intelligent nodes will take all necessary decisions in order to support the required functionality (including guaranteeing the QoS). Thus, the Service/Network infrastructure will:

- Deal with active content objects, as well as unstructured bitstreams
- Provide the required transport, congestion-control, policy, signalling and processing protocols in a distributed manner
- Provide all necessary Fault, Configuration, Accounting, Performance and Security (FCAPS) functionality
- Support mobility and portability of content, information, services, terminals and users
- Support virtualisation and self-organisation/self-management
- Support robustness, stability, virtualisation and survivability
- Hide unnecessary complexity and provide information/resources hooks
- Monitor and control network resources
- Offer hooks for supporting the Quality of Service (QoS) and Service Level Agreements (SLA).
- Offer connectivity to multiple administration and cross-network domains (including sensor networks, mobile networks, corporate networks)
- Offer security and privacy of the infrastructure
- Offer interworking and Interoperation capabilities

The medium layer is the **Distributed Content/Services Aware Overlay**. Content-Aware Network Nodes (e.g. core routers, edge routers, home gateways, terminal devices) will be located at this overlay. These nodes will have the intelligence to filter the content and Web services that flow through them (e.g. via Deep Packet Inspection, DPI) or identify streaming sessions and traffic (e.g. via signalling analysis). Alternatively, the content may be considered formulating *information objects* as first order elements in the network, thus be directly identifiable by the network nodes⁷. In either case, the Nodes of this group will recognise and qualify the content. Part of this information may be stored locally and/or reported to the higher layer of hierarchy (Information Overlay).

⁷ 4WARD project. Deliverable D-6.1: "First NetInf architecture description"

Content/Services Aware Overlay may be split in multiple virtual overlays/clouds (Figure 6), which may be dynamically constructed at the layers between the content and the information overlays. We may consider overlays for content caching, content classification (even content indexing in the future), network monitoring, content adaptation, optimal delivery/streaming. With respect to content delivery, nodes at this layer may operate as hybrid client-server, peer-to-peer or cloud networks, according to the delivery requirements.

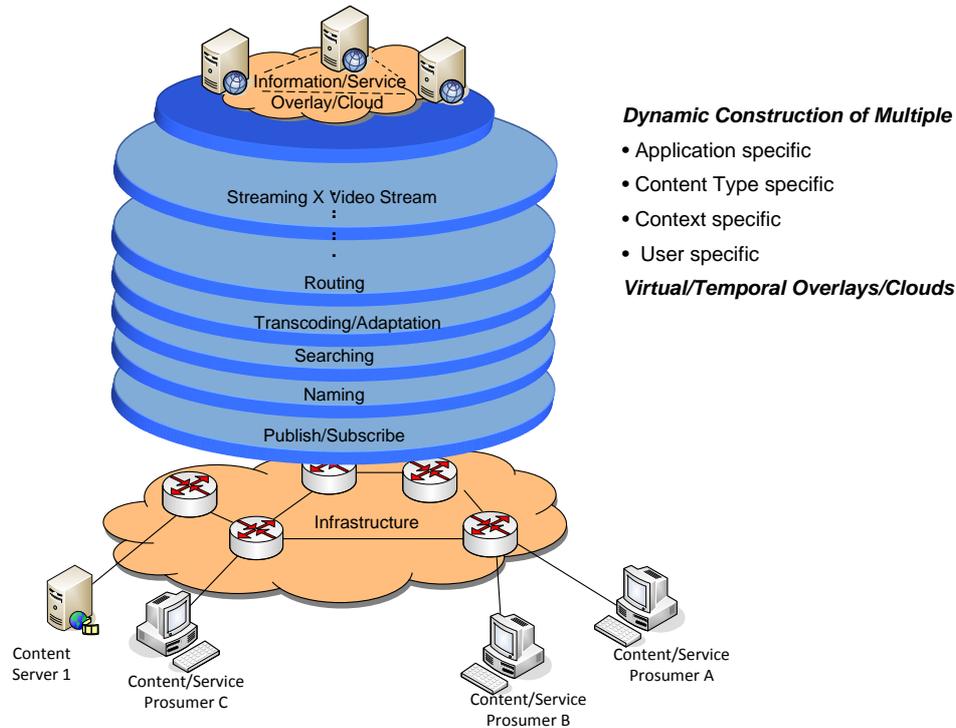


Figure 6: Dynamic construction of Virtual/Temporal Overlays/Groups

In summary the Distributed Content/Services Aware Overlay will:

- Provide content (media and services) awareness (inspection, crawling, recognition, categorisation, indexing)
- Provide content caching and “findability”
- Provide content (media and services) adaptation and personalisation
- Offer mechanisms for content (media and services) extraction, combination, creation, orchestration
- Offer mechanisms for name resolution, route by name and route by type capabilities
- Offer mechanisms for accountability and billing of content (media and services)
- Offer mechanisms for protection of content (including beyond DRM, lightweight management of the content)

At a higher layer, it is the **Information Overlay**. It will consist of intelligent nodes or servers that have a distributed knowledge of both the content/web-service location/ caching and the (mobile) network instantiation/conditions. Based on the actual network deployment and instantiation, the service scenario, the service requirements and the service quality agreements, these nodes may vary from unreliable peers in a next-P2P topology to secure

corporate routers or even data centres in distributed carrier-grade cloud networks. The content may be stored/ cached at the *Information Overlay* or at lower hierarchy layers, though the *Information Overlay* will be always aware of the content/services location/caching and the network information. Based on this information, it may decide on the way that content will be optimally retrieved and delivered to the subscribers or inquiring users or services.

In summary the Information Overlay will:

- Offer mechanisms for control the content and service “findability”, construction, orchestration, representation and deployment.
- Offer mechanisms for supervise the content/services location/caching/deployment and the network resources information.
- Offer mechanisms for content (media and services) publishing/pushing/subscription
- Offer mechanisms for accountability and billing of content (media and services)
- Offer mechanisms to support networking tussle and new business models

Finally, at the highest layer the ***Application’s layer*** is located. Applications will use efficiently the services, the information and the media/content provided by the content-centric architecture and offer novel media experiences to the users.

4.2 Autonomous Content Centric Architecture

In the previous section, we assumed that the network may identify the content using a number of methods (e.g. Deep Packet inspection, session signalling). In a more forward looking approach, moving from an evolutionary to a more clean-slate approach, we may introduce the concepts of “**autonomous content-centric network architecture**” based on the “**Content Objects**”.

In Figure 2, we have already introduced the concept of content components, which can ad-hoc, on the fly generate/reconstruct semantically enriched 3D augmented/ virtual worlds in order to create an orchestrated immersive media experience. Here we further expand this concept to “**Content Objects**”.

An **Content Object** (or simply object) is a polymorphic/holistic container, which may consist of media, rules, behaviour, relations and characteristics or any combination of the above.

- **Media** can be anything that a human can perceive/experience with his/her senses (a dancing person, the second violin in a symphonic performance, a tear on your cheek).
- **Rules** can refer to the way an object is treated and manipulated by other objects or the environment (discovered, retrieved, casted, adapted, delivered, transformed, presented). Rules can for instance be used to specify if the media in the object would allow rescaling and that it would accept a delivery delay of 2 seconds, but that it should certainly arrive for presentation at the end-user side before a child object: the object knows its purpose in the integrated media experience and therefore its priority for transfer. Also the options for manipulation by the end-user at the moment of presentation could be included.

- **Behaviour** can refer to the way the object affects other objects or the environment.
- **Relations** between an object with other objects can refer to time, space, synchronisation issues. Relations could for instance describe that an audio object of a singing person is related to an animated 3D model of the singer and that lip synchronisation is required.
- **Characteristics** meaningfully describe the object and allow to retrieve related objects: user interaction with a ‘coq-au-vin’ object may visualise in the immersive 3D environment the ingredients and their current prices or may lead to the ad-hoc building of 3D replicas of the restaurants where the dish is available.

Objects can be **hierarchically organised**, like the constituting instrument channels of a music band, and can trigger the generation of new objects. An object can be **divided/ split** into new objects or multiple objects can be **combined/merged** and finally create new objects, and these operations may happen while travelling over the network.

An object can be **cloned**. The clone keeps the characteristics of its “parent” object but knows that it is a clone. This is also associated with issues like caching (object lifetime, check for updates) and Digital Rights Management (DRM).

The autonomous objects will travel over the network, split and combined to generate the new service or virtual world object. The Future Content Centric Internet will support the content objects in order to meet their relations.

It is currently very difficult to imagine what a network architecture that support objects would look like. An attempt to map the characteristics of the layered approach which is depicted in Figure 3 into the novel “layer less” concept of the object is shown in the Figure 7, where one or more layers are mapped to one or more characteristics of the object.

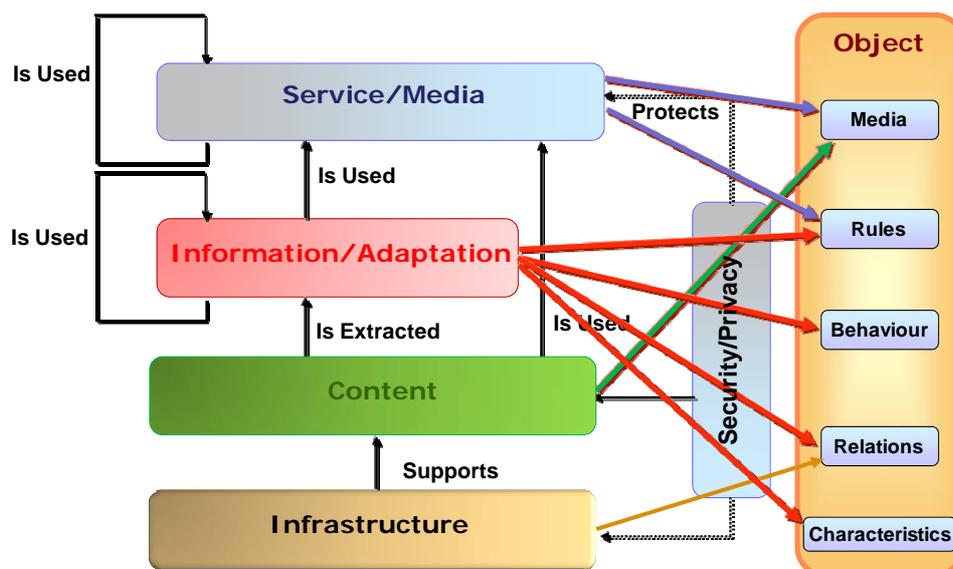


Figure 7: Mapping a layers-based Content-Centric Internet Architecture into “Objects”

More specifically, transfer and integration of objects for the purpose of the creation of an orchestrated “Media” experience clearly demands intelligence that combines application



(“Service/Media”) and “Content” information. The intelligence could be embedded in the objects themselves, retrieving information from the network and providing instructions for routing and transformation, or the intelligence could be hosted in network nodes that attempt to satisfy the requests of the objects as they are described in the “Rules”, “Behaviours” and “Relationships” (which take input from the “Information/Adaptation”, “Content” and “Infrastructure” layers) . Finally, the “Characteristics” that meaningfully describe an object take, mainly, input from the “Information/Adaptation” layer. It should be noted that both the Logical and the Autonomous Content-Centric Internet Architectures provide a holistic view of our vision for a FCCIA, which apart from serving the demanding real-time applications of the future offer also the possibility of reducing significantly the network management barriers and limitations of the existing approaches

5 Conclusions

The ultimate goal of this document is to initiate a debate between all the interested stakeholders with respect to potential Future Content-Centric Internet Architectures. Initially, we analysed the reasons on why a Content-Centric approach will be one of the main pillars of the Future Internet and the limitations and barriers of the current Internet. Then, the most important “content-centric” design requirements and principles of the Future Content-Centric Internet Architecture were reported, mainly focusing on enhancing the media experience of future users.

Our aim is to start formalising the requirements and principles that will lead to a concrete Future Internet Architecture, as it is envisioned from the Future Content Networks perspective. In this attempt, it was considered that the Future Internet Architecture requires a holistic design approach, which spans from network infrastructure up to content and services, which in turn may affect the design of the network infrastructure.

Summing up, we believe that in order to achieve the vision of a Future Internet fully suited to future users’ needs, several aspects need to be considered. Among others, network structure complexity vs. engineering design simplicity, scaling vs. delivering quality and response time, efficiency vs. user friendliness, services and content location, user and network mobility, societal aspects and issues of trust and security, just to name a few. Moreover, the decision on following a revolutionary or a clean-slate approach is heavily under discussion.

6 Acknowledgements

This work is based on discussions carried out by a selected group of experts from the Future Content Networks (FCN) session of the Future Internet Assembly (FIA) and expresses their opinion. Yet, the authors would like to acknowledge input from a number of documents mainly generated from the Future Media and 3D Internet Task Force (FM3DI-TF), the Media Delivery Platforms (MDP) and the User-Centric Media (UCM) clusters of projects and other selected projects, RFCs and studies have been taken into account.

7 Annex – Relative FCN fora and groups

It is important to understand that any architectural evolution towards Future Internet will not be successful if it is proposed in isolation. Today, several frameworks of industry recommendations and standards exist, established to promote agreement and interoperability between many of the organisations who will ultimately be the major commercial players of the Future Internet. This section describes some of the most relevant frameworks and the extent to which they are already addressing content-centric architectures.

One key property of these frameworks is that a) they provide guidelines on the application of the same core standards, but in the context of specific use cases which their industry contributors believe will be commercially relevant for future consumers, and b) mechanisms to develop new use cases as the Future Internet evolves.

7.1 Digital Living Network Alliance (DLNA)

The Digital Living Network Alliance⁸ is a cross-industry organisation of consumer electronics, computing industry and mobile device companies. It has a vision of seamless interoperability between consumer electronics, personal computers and mobile devices both inside and outside the home. Its primary goal is to promote the adoption and sharing of new digital media and content services. The DLNA is developing interoperability guidelines based on open industry standards to complete the cross-industry digital convergence, and there are already a range of mass-market products which have been certified to meet the guidelines published so far.

In the context of the Future Content-Centric Internet, DLNA has defined a set of categories for networked devices (device classes) and a subset of existing standards at the Application Layer which these devices must support in order to become DLNA certified. The DLNA guidelines provide a comprehensive model for the ecosystem of content-centric devices which will surround users of the Future Internet. For example, they define the capabilities of content-aware nodes such as digital media servers, players and renderers and the signalling and media transfer protocols between them. They address issues of networking and connectivity, device discovery and control, content formats and distribution, and DRM and content protection.

7.2 IP Multimedia Subsystem (IMS)

IMS (IP Multimedia Subsystem) architecture provides the infrastructure for high quality multimedia services in environment totally controlled by the network operators. This tight control and network operator centric approach is very different from the key principles of Future Internet and freedom of choice, openness and lack of barriers to innovation.

However, there are some characteristics of the IMS architecture that can be adopted in the design principles of the Future Internet like the use of an open and standard signalling protocol (SIP) that allow the integration of new services. The way new services are developed in IMS, sharing all of them the common functionalities provided by the IMS infrastructure, is a key approach for the successful of Content-Centric Future Internet services. The characteristics of SIP signalling protocol, the most widely used signalling protocol in current Internet, let us

⁸ http://www.dlna.org/news/DLNA_white_paper.pdf



see that in the future designs and development, the use of this or new standard and open signalling protocols will be a must in the Future Internet.

The lesson learnt in the current use of IMS frameworks is that the trend is not the use a network operator tight control framework, but the use of a common and open signalling protocol that could be the base and not the barrier for the development of new multimedia services.

7.3 Open IPTV Forum

The Open IPTV Forum is a pan-industry initiative with the purpose of producing end-to-end specifications that will take the next generation of IPTV to the mass market⁹. The Forum recognises that the IPTV value chain may differ across geographies and markets, and thus supports the existence of multiple entities with different regions of administrative control and ownership interests. The Forum has defined an Open IPTV Architecture which includes the following domains: Consumer, Network Provider, Platform Provider, IPTV Service Provider and Content Provider. Being an end-to-end specification, a significant proportion of this architecture overlaps with the logical content-centric architecture described above.

The Forum's specifications are comprehensive and take into account a wide range of TV-centric applications and services which could be offered over the Future Internet beyond just streamed television and content on demand. They make provision for Application and IMS Gateways to enable interoperation between IPTV and communications services, as well as interworking with DLNA devices. They also address service discovery and a range of different content delivery scenarios, including the concept of clustered sets of delivery nodes.

7.4 Conclusion

A Future Content-Centric Internet Architecture will not be accepted if it proposes a disconnect with the roadmaps already proposed by organisations which have significant support from key industry players, many of whom will have a major role in its implementation and management. Following the proposal of the Architecture at a high level, further work is required to identify in detail where existing specifications can be adopted and promoted as the Architecture moves forward.

⁹ <http://www.openiptvforum.org/aboutus.html>

8 List of Contributors

8.1 Editorial Team

Petros Daras	CERTH/ITI	daras@iti.gr
Doug Williams	BT	doug.williams@bt.com
Carmen Guerrero	UC3M	guerrero@it.uc3m.es
Ian Kegel	BT	ian.c.kegel@bt.com
Isidro Laso	European Commission	Isidro.LASO@ec.europa.eu
Jan Bouwen	Alcatel-Lucent	jan.bouwen@alcatel-lucent.be
Jean-Dominique Meunier	Thomson	Jean-Dominique.Meunier@thomson.net
Norbert Niebert	Ericsson	norbert.niebert@ericsson.com
Theodore Zahariadis ¹⁰	Synelixis Ltd	zahariad@synelixis.com

8.2 List of Contributors

The Future Content Networks (FCN) group includes the following members, who have contributed either directly or indirectly in the formation of the ideas that are included in this document via emails and comments.

Name	Company	Project
Aiko Pras	University of Twente	Emanics
Alex Conconi	TXT	PlayMancer
Andrew Oliphant	BBC	TA2
Borje Ohlman	Ericsson	
Carmen Guerrero	UC3M	Content
Christian Timmerer	Universität Klagenfurt	
Cornelius Hellge	Fraunhofer HHI	
Dimitri Papadimitriou	Alcatel-Lucent	
Doug Williams	BT	TA2
Ebroul Izquierdo	Queen Mary University	PetaMedia
Eugenia Fuenmayor	Barcelona Media	

¹⁰ Document Editor

Federico Alvarez	UPM	SEA
Fidel Liberal	University Basque Country	Adamantium
Guillermo Cisneros	UPM	SEA
Harilaos Koumaras	NCSR Demokritos	Adamantium
Ian Kegel	BT	TA2
J. C. Point	JCP-Consult	Mobithin
Jan Bouwen	Alcatel-Lucent	UserMedia
Jari Ahola	VTT	P2P Next
Jean-Dominique Meunier	Thomson	4NEM
Jean-Dominique Meunier	Thomson	4NEM
Jean-Noel Patillon	Motorola	
Emilio Leonardi	Polito	NapaWine
Luca Celetto	STM	SEA
Norbert Niebert	Ericsson	4WARD
Olivier Festor	LORIA	
Peter Stollenmayer	Eurescom	TA2
Petros Daras	CERTH/ITI	Victory
Pierre Yves Danet	Orange/France Telecom	4NEM
Robert-Paul Berretty	Philips	Content
Sergio Gusmeroli	TXT	COIN
Stefan Debal	Microsoft	
Theodore Zahariadis	Synelixis	SEA
Thomas Schierl	Fraunhofer HHI	SEA
Wolfgang Nejd	L3S Research Center	Pharos
Yves-Marie Le-Pannerer	Thomson	4NEM