

Towards a Content-Centric Internet

Theodore ZAHARIADIS^{a,1}, Petros DARAS^b, Jan BOUWEN^c, Norbert NIEBERT^d,

David GRIFFIN^e, Federico ALVAREZ^f, Gonzalo CAMARILLO^g

^aSynelixis Solutions Ltd, 10 Farmakidou Av., Chalkida, Greece

^bCERH/ITI, Themi, Thessaloniki, Greece

^cAlcatel-Lucent, Copernicuslaan 50, Antwerp, Belgium

^dEricsson, Kackertstraße 7-9 52072 Aachen, Germany

^eUniversity College London, Gower Street, London, UK

^fUniversidad Politécnica de Madrid, Madrid, Spain

^gEricsson, Hirsalantie 11, Jorvas 02420, Finland

Abstract. In most cases, current Internet architecture treats content and services simply as bits of data transported between end-systems. While this relatively simple model of operation had clear benefits when users interacted with well-known servers, the recent evolution of the way the Internet is used makes it necessary to create a new model of interaction between entities representing content. In this paper we study the limitations of current Internet and propose a new model, where the smallest addressable unit is a content object, regardless of its location.

Keywords. Content-Centric Internet, Content Networks, Internet Architecture

Introduction

The Internet as we know it today is heavily based on a model that interconnects interfaces of end-hosts – both servers and user devices – that are usually identified by IP addresses. The wealth of information and applications we enjoy today is all hosted on computers and invisible to the basic operation of the Internet, which treats content and services simply as bits of data transported between end systems. While this relatively simple model of operation had clear benefits in the early days of the Internet when users interacted with well-known servers using services such as file transfer or remote terminal access, the recent evolution of the way the Internet is used makes it necessary to create a higher level platform for the interaction with digital entities representing content of all kinds – to make a Content-Centric Internet (CCI) rather than today's computer-centric one. In a CCI the content is addressable, regardless of its location. In the following, we will study current Internet's limitations, analyse the benefits of a CCI approach, and provide the design principles and the requirements of a Content Centric Internet architecture.

¹ Corresponding Author: Theodore Zahariadis, Synelixis Solutions Ltd, 10 Farmakidou Av. GR 34100, Chalkida, Greece; Email: zahariad@synelixis.com

1. Limitations of the Current Internet Architecture

Today, the vast majority of the Internet usage concerns content and services discovery & retrieval, content delivery and streaming and Web services access. The user cares only about the content or service itself and proper delivery, while he/she is oblivious to their location. That is, the user knows that he/she wants news from the BBC, videos from YouTube or weather information, concrete and delivered in suitable quality and format, but does not know or cares on which machine the desired data or service resides, as soon as reliability, security and privacy are guaranteed. This functionality is realised by the current Internet Architecture as shown in Figure 1. It consists of the following types of nodes:

- a) **Content Servers or Caches**,
- b) Centralised, decentralised or clustered **Servers**, including *Search Engines* and *Supporting Servers* (e.g. DNS servers, AAA servers, DRM servers, etc.),
- c) Core and edge **Routers** and **Residential Gateways** (represented as R1 to R5) and
- d) Users, connected via fixed, wireless or mobile **terminals**.




Figure 1. 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 or inspect the routed packets 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, she 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 1 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). 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? When users demand image resolutions in video that require bandwidths greater than can be supported over

the typical link lengths? If 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. However some changes to this schema would make better use of the available resources. For example:

- a) if ***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,
- b) if ***routers could identify/analyse what content is flowing through them***, the search engines would gain much better knowledge of (even the streaming) content location and provide information even on “live” video streams,
- c) if the ***network could identify what is the best path to the user*** (less congestion, lower delay, more bandwidth), it could provide a better way to deliver data
- d) if ***content could be selected and adapted to the context***, the user would have a much easier life e.g. when entering a living room, a phone TV session could transfer to the big screen and adapt to the resolution offered there.

2. From Content to Services to Media Experiences

In the debate about the shape of future Internet, three powerful concepts drift to the surface, vying for attention: User, Service and Content. Each of the three presents itself as a powerful force that is able to explain recent evolution and that claims the right to drive the future Internet. The ***user-centric*** perspective emphasises the end-user experience as the driving force for all technological innovation, observing how today the Internet is a network of active social users rather than a connection of devices. The ***service-centric*** view has roots in both enterprise IT solutions and the Web 2.0 mash-up culture, showing how valuable applications can be built faster and more efficiently if service components can be reused in flexible ways. The ***content-centric*** view refers to the central role that rich media content is playing in attracting users to Internet services, as content consumers are increasingly also as content producers, and how the transfer of media content can impact the network operation. As the three views are emphasising different aspects rather than expressing opposing statements, merging or homogenizing towards an encompassing perspective may help towards the right design choices for a future Internet.

To satisfy user experience, a content-centric Internet will depend on the realisation of a set of content-specific services. In this way, the content-centric perspective adds new service components to the service-centric view. Content-centric services include content distribution networking for both on-demand and live media distribution, content publishing, discovery, adaptation and processing services, DRM services, conferencing services, media annotation, indexing and search services. Figure 2 shows the interrelations between the different components. In more details, we may define:

- ***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.

- **Content** is any type and volume of raw information that can be combined, mixed or aggregated to generate new content and media. Content may be pre-recorded, cached or live, static or dynamic, monolithic or modular.
- **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.
- **Security and Privacy** will be a property of *content*, *information*, *services* and Infrastructure, allowing much more efficient control over *content* objects.
- **User/Media experience** encompasses all aspects of the end-user's interaction with the services and the Media. True user experience goes far beyond giving customers what they say they want, or providing checklist features.




Figure 2. Future Content-Centric Internet components interrelation

2.1. Impact of the User-Centric Perspective

Taking the end-user with his/her needs and desires as the initiating force for the design of the Future Internet and the applications it will support, we consider the following requirements that will have an impact on the **Service/Media** layer:

- **The end-user is the endpoint**, rather than his/her device. Users should be able to easily find each other and engage in interactions, even if they use multiple devices in parallel without falling in the trap of limiting the user to one single identity.
- **Universal accessibility for services** with user experience; various users will approach the offered services with different levels of competency, and this level will evolve over time as they make use of the service.
- **Universal accessibility for content generation**; users already act as content producers, implying that if a Service-Centric network offers components for end-user experience creation, they should be usable by all users.
- If the **network is enhanced with a certain intelligence** to optimize the user experience, this should not lead to a feeling of loss of control with the user, dealing with an unpredictable environment.

2.2. Impact of the Service-Centric Perspective

Although the Internet is supporting a wide variety of applications, several functional building blocks are common between large groups of applications. The Service-Centric perspective therefore argues that it makes sense to design a network environment that supports the flexible creation, publishing, discovery and use of common service components. Flexibility here refers to easy location-independent detection and invocation of service components. Just-in-time inclusion of service components – i.e. at the moment of the creation of the end-user experience – allows optimization of network services and supports rapid innovation. The common service components that are typically listed include user identification, authentication and authorization, security and DRM, bandwidth management, storage, power management, payment, location and time context information, user activity, content adaptation, search and indexing functions. Some of these – user authentication and authorization, content adaptation to devices and user context – can be seen as driven by the User-Centric perspective.

2.3. Impact of the Content-Centric Perspective

The Content-Centric perspective highlights the driving role that rich multimedia content has played and is expected to play in the growth of the Internet, in terms of usage and traffic. The web has become a true Media Web, and the volume of transferred content will continue to rise sharply, as the quality of the media content further increases (High-Definition and Ultra High-Definition Content, 3D and stereoscopic content, multi-view content etc), as more experience of content becomes active and social and as more users evolve from mere consumers to active creators and/or repurposes of content. The Media Web is furthermore evolving to a Real-Time Media Web with live content streams and multimedia person-to-person or group communication. This can be a separate application experience or it can be embedded in frame experiences like gaming, education and users collaboration.




Figure 3. The convergence of three different perspectives

To satisfy user experience, a Content-Centric Internet will depend on the realisation of a set of content-specific network services. In this way, the Content-Centric perspective adds new service components to the Service-Centric view - indispensable for media experiences, or emphasises already identified ones. Content-Centric services include content distribution networking for both on-demand and live

media distribution, content publishing, discovery, adaptation and processing services, DRM services, conferencing services, media annotation, indexing and search services. Figure 3 schematically depicts the convergence of the three different, yet complementary perspectives.

3. The Concept of Content Objects

Currently, media content is the result of an off-line, cumbersome and lengthy creation process, whereby content components are composed into a meaningful and appealing presentation. The distribution over the network for consumption is then the transfer of the finalised complete media presentation in the form of bit streams, followed by a play-out at the end-user's device. The key concepts for the Service-Centric perspective as explained above are the identification and separation of meaningful service components and the just-in-time on-the-fly flexible integration of such components into an application experience. It is expected that this evolution for software and network functions will also take place for rich media, i.e. that media experiences will be created as the just-in-time composition of content component objects that are easily located, synchronised, reused and composed.

Such an approach can already be discerned in virtual world applications where users contribute to the content creation: the virtual world representation on the end-user's device is the composition of objects that have been created by various authors and are fetched as they are required for representation. Figure 4 represents a possible mixed-reality scene for a person-to-person interaction, combining stored and live media objects from a multitude of sources and engagement of many senses.




Figure 4. On the fly generation/reconstruction of semantically enriched worlds

The availability of the constituent content objects 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 collaborative audio-visual content creation.

- 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 greatly facilitated.
- The combination of captured audio-visual content with synthetic 3D content creates exciting mixed-reality experiences.
- The possibility for the user to actively intervene and mould content components through natural, non-verbal interfaces enhances the experience and provides authoring capabilities to the users, e.g. by enabling them to reshape, personalize, and re-experience in unique ways audiovisual content, including layered metadata.

3.1. A forward-looking alternative for content objects

The classic layered approach may not be the ideal match for the content object vision: the advanced content treatment service functions that are required may exhibit characteristics that differ substantially from the non-content-driven service components, leading to the definition of service components that are positioned in a blurred area between content, service and user layers. An alternative is a clean-slate approach for the network design, starting from the content object itself, a content-centric network architecture, the **Autonomic Layer-Less Object Architecture (ALLOA)**. In Figure 4, we have already introduced the concept of content objects, 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*”. A *Content Object* is an *autonomous, polymorphic/holistic container*, which may consist of media, rules, behaviour, relations and characteristics or any combination of the above.

- **Media** are the actual content pixels. It 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, and 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, and 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 retrieval of its 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/spit* into new objects or multiple objects can be *combined/merged* and finally create new objects, and these operations may happen while objects are “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 cloning has implications in the opening of novel business models around “actively experience audiovisual content”: for example, the same audiovisual content can be distributed with different characteristics: for example for a music, ranging from a simple MP3 file to a more complex music content in which the user has authoring capabilities to reshape the music piece, thanks to the availability of further metadata and higher level representations enabling this user a number of degrees of freedom in real-time, e.g. in terms of re-orchestrating the music, re-arranging (post-production), shared (social) active fruition.

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

An example of a content object is shown in Figure 5. We can assume that “Barbie” is a Content Object mash-up. It consists of different other basic objects E.g. It *has* “skin”. The “*skin*” is a content object, which has different colours or textures. It *has* hair. The “*hair*” is a content-object, which has different colour, length, style. It *has* “eyes”. The “*eyes*” is a content-object, which has different colour, length, style, shadows. It *wears* “cloths” and *carries* “accessories”. Both “cloths” and “accessories” are components mash-ups, which may change based on the time, context (school, dance, gym), emotions. What really differentiates this example from today’s technologies is that any digital object may be decomposed or composed into Content Objects: the environment, the cars, the buildings, the things, etc., generating in real-time new content in the same sense that SOA may support service mash-ups.



Figure 5. A Barbie Content Object

Another potential use of the Content Object approach can be seen in Figure 6, where a user simply sketches (in two-dimensions) the scene of interest. Each individual item of the sketch can be thought as a Content Object, which contains all the information that characterises it. Thus, each Content Object can serve as an *autonomous rich item*, containing both low-level and high-level features. Similar 3D objects (to each 2D sketched item) are then automatically retrieved from the CCN and placed in a 3D environment which can also be a Content Object.




Figure 6. From user generated 2D sketches to 3D worlds

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 2 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 entities of the object.

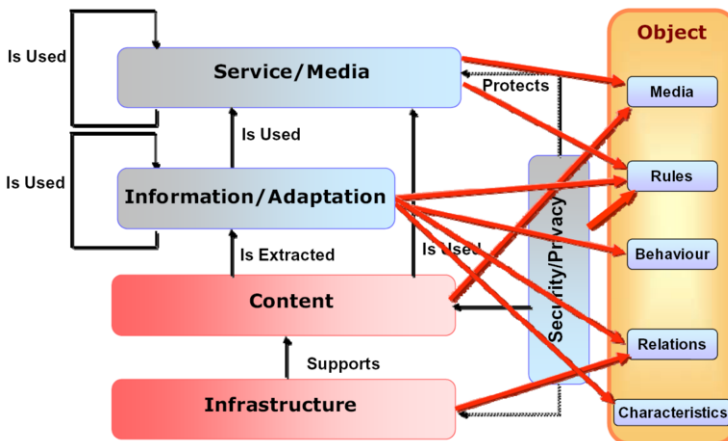


Figure 7. Mapping a layer-based Content-Centric Internet Architecture into “Objects”

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.

4. Conclusions

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. Yet, an incremental approach starting from a Virtualised Network towards Content Object Mash-ups is a possible scenario.

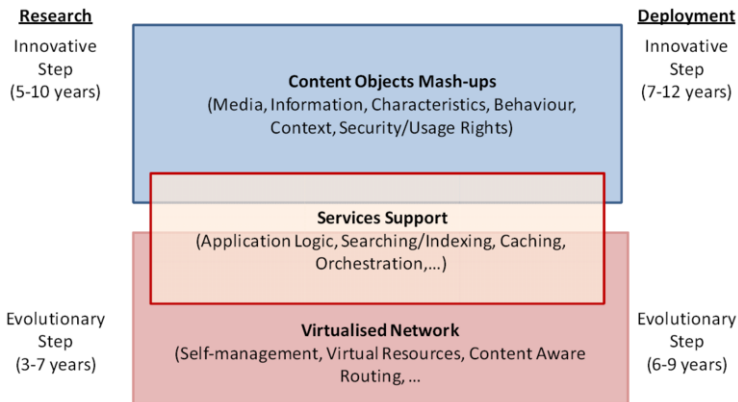


Figure 8. Research and Deployment for Future Content Centric Internet

5. Acknowledgements

This work is based on discussions carried out by a selected group of experts from the Future Content Networks session of the Future Internet Assembly (FIA) and expresses authors' opinion. Yet, the authors would like to acknowledge input from a number of documents mainly generated from the Future Media Internet (FMI) Task Force, the Media Delivery Platforms (MDP) and the User-Centric Media (UCM) clusters and selected projects, RFCs and other studies. Part of this paper has been discussed in the EU funded projects COAST, ENVISION, I-Search and nextMEDIA.

References

- [1] R. Bush, D. Meyer, "Some Internet Architectural Guidelines and Philosophy," RFC 3439, Dec. 2002
- [2] W. Willinger, J. Doyle, "Robustness and the Internet: Design and evolution", 2002, http://www.maoz.com/~dmm/complexity_and_the_internet/robustness_internet_design_evolution.pdf
- [3] 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
- [4] D. Clark, J. Wroclawski, K. Sollins, R. and Braden: *Tussle in Cyberspace: Defining Tomorrow's Internet*, 2002 Conference on Applications, Technologies, Architectures, and Protocols For Computer Communications, SIGCOMM '02, (2002), New York, NY, 347-356
- [5] Future Content Networks Group, "Why do we need a Content-Centric Internet? Proposals towards Content-Centric Internet Architectures," White paper. 2009