“Future Media Internet Architecture Think Tank”

White Paper

Future Media Internet Architecture Reference Model (v1.0)

Acronym	FMIA-TT
Group URL	www.fi-nextmedia.eu / http://www.gatv.ssr.upm.es/nextmedia/index.php?option=com_content&view=article&id=34&Itemid=40
Document Name	FMIA-TT reference model outline
Abstract (for dissemination)	This document is the pre-final version of a reference model for the Future Media Internet Architecture.
Keywords	Future Media Internet, Architecture,
Date	1 March 2011
Status	Final
Version	1.0

Table of Contents

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

The Future Internet (FI) is expected to be a communication and delivery ecosystem, which will interface, interconnect, integrate and expand today's Internet, public and private intranets and communication networks of any type and scale, in order to provide efficiently, transparently and securely highly demanding services to humans and systems.

This complex networking environment may be considered from various interrelated perspectives: the networks & infrastructure viewpoint, the services viewpoint, the media & information viewpoint.

This document has been produces by a discussion forum of experts in the area of media and networks, named: Future Media Internet - Think Tank (FMIA-TT). FMIA-TT aims to create a reference model of a "Future Media Internet Architecture", covering delivery, in the network adaptation/enrichment and consumption of media over the Future Internet ecosystem. This white paper concludes the first phase of the working group by proposing a FMIA reference model.

1 Introduction

Today, Internet is the most important information exchange ecosystem. It has become the core communication environment not only for business relations, but also for social and human interaction. The immense success of Internet has created even higher expectations for new applications and services, which the current Internet may not be able to support. Advances in video capturing and encoding have lead to massive creation of new multimedia content and applications, providing richer immersive experiences: 3D videos, interactive environments, network gaming, virtual worlds, etc. Thus, scientists and researchers from companies and research institutes world-wide are working towards realising the Future Internet.

The Future Internet (FI) is expected to be a holistic information exchange ecosystem, which will interface, interconnect, integrate and expand today's Internet, public and private intranets and networks of any type and scale, in order to provide efficiently, transparently, timely and securely any type of service (from best effort information retrieval to highly-demanding, performance critical services) to humans and systems.

This complex networking environment may be considered from various interrelated perspectives: the networks & infrastructure perspective, the services perspective and the media & information perspective. ***The Future Media Internet is the FI viewpoint that covers the delivery, in-the-network adaptation/enrichment and consumption of media over the Future Internet ecosystem.***

Significant efforts world-wide have already been devoted to define, build and validate the FI and/or some of its pillars: the NetSE [1] programs as successors of FIND [3] and GENI [4] in USA, the AKARI [5] program in Japan, the Future Internet [6] program in Korea. In Europe, a significant part of the FP7 has been devoted to the FI [7]; both large and small/targeted research projects have already reported some preliminary results, others are in progress (4WARD, ENVISION, COAST, OCEAN, NEXOF-RA, IoT-A, Nanodatacenters), while the recently announced Private Public Partnership (PPP) program is expected to provide an industry-driven, holistic approach encompassing R&D on network and communication infrastructures, devices, software, service and media technologies.

In this rapidly evolving environment, the aim of the Future Media Internet - Think Tank (FMIA-TT) is to build a discussion forum for experts in the area of media and networks, which will brainstorm, analyze, debate, agree and propose a reference model of a "Future Media Internet Architecture", covering the delivery, ***in-the-network adaptation/enrichment*** and consumption of media over the Future Internet ecosystem. The overall editing and homogenization has been provided by the nextMedia CSA project consortium.

2 Current Internet content delivery limitations

We start this section by reviewing how content discovery, retrieval and delivery take place in the current Internet. Users want news from CNN, videos from YouTube or weather information, but they do not know or care on which machine the desired data or service reside. Information/content retrieval and delivery may be realised by todays Internet network architecture as shown in Figure 1. The network consists of: a) *Content Servers* or *Content Caches* (either professional or user generated content and services), b) centralised or clustered *Search Engines*, c) core and edge *Routers* and optionally

Residential Gateways (represented as R1 to R5) and d) Users connected via fixed, wireless or mobile terminals.

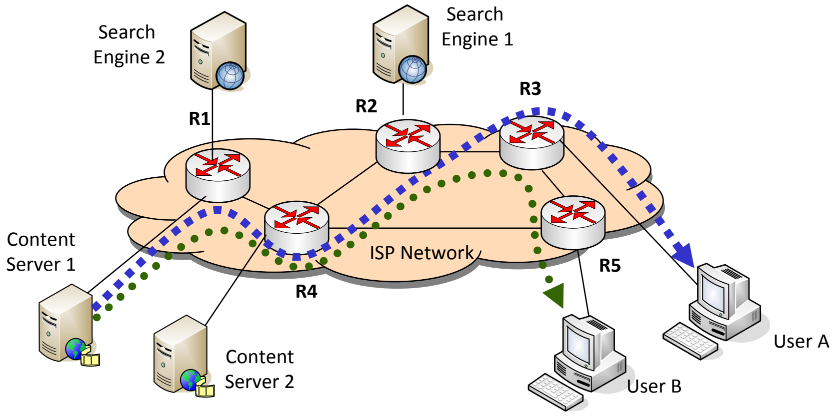

Figure 1: Today's Internet Architecture

The first step is *Content Discovery by the Search Engines*: the Search Engines crawl the Internet to find, classify and index content and/or services. The second step is *Content Discovery by the User*: the user queries a Search Engine and gets as feedback a list 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 him.

In order to show with an example the limitations of today's Internet, let us consider the simple case of the delivery of a popular video from Content Server 1 in Figure 1 (e.g. a YouTube server). If a few dozen of users from a large building block request a video, the same video chunks will be streamed a few dozen of times. If a neighbourhood has a few dozen of blocks, and a city a few hundreds neighbourhoods, the very same video chucks will traverse the same network links thousands of times. If we continue aggregating at country and world-wide level, we will soon run out of existing bandwidth just for a single popular video stream.

This means that the three steps of content discovery and delivery can be significantly improved:

- ***(In the network) Dynamic caching:*** If the *content could be stored/cached closer to the end users*, not only at the end-points as local proxies but also transparently in the network (routers, servers, nodes, data centres), then content delivery would have been more efficient.
- ***Content Identification:*** If the *routers could identify/analyse what content is flowing through them*, and in some cases are able to *replicate* it efficiently, the search engines would gain much better knowledge of the content popularity and provide information -even when dealing with "live" video streams.
- ***Network topology & traffic:*** If the *network topology and the network traffic per link were known, the best end-to-end path* (less congestion, lower delay, more bandwidth) would be selected for data delivery.
- ***Content Centric Delivery:*** If the *content caching location*, the network topology and traffic were known, more efficient content-aware delivery could be achieved based on the content name, rather than where the content is initially located.
- ***Dynamic Content Adaptation & Enrichment:*** If the *content could be interactively adapted and even enriched in the network*, the user experience would be improved.

3 Design principles for a Future Media Internet architecture

In this section social, economic and policy forces rather than technological aspects are considered for the extraction of the FMI design principles[1]. From a content-centric viewpoint, three main principles emerge and are highlighted next.

3.1 KISP (but not too simple)

The KISP (Keep It as Simple as Possible) principle is based on the famous quote by Albert Einstein: "Make everything as simple as possible, but not simpler". Complex problems sometimes require complex solutions and the FI will be providing non-trivial functionality in many respects. However, designers should keep in mind this principle and prefer relatively simpler and more elegant solutions than over-engineered designs. Complex systems are generally more difficult to manage and less reliable since more things may malfunction at any given time. Therefore, complexity should always be added for a good reason.

3.2 Design for Tussle

This design principle states that the Future Internet should not be engineered to favour one particular Internet stakeholder over another. The FI should be capable of supporting flexible business models where multiple stakeholders can participate in an open environment that supports and encourages innovation and participation without barriers. Open architectures and protocols will enable increased competition between providers (including network, service and application providers) increasing quality and value to the benefit of all. Individuals should be able to produce as well as consume content; innovators, both small and large, should be able to introduce new products, new technologies and even new communication paradigms without the hindrance of conformity to established or traditional business models [15]. The FI should support a greater participation of individuals, communities and small businesses alongside larger and more established organizations and the FI should enable all providers of content, services or other forms of added value to receive appropriate compensation for their contribution.

3.3 Sustainability

The FI must be designed as a sustainable network being flexible enough to continuously evolve, develop and extend in response to changing societal requirements. Adopting such a sustainable design will allow for environmental and societal developments over many decades, making the FI able to support universal communication that will overcome the obstacles of language, culture, distance, or physical ability which exist in the current Internet (CI). The sustainability of the FI will rely on its ability to be scalable, available and reliable in a resource- and cost efficient manner. The latter means that the 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). Finally, the FI should be able to provide openness to users to facilitate the

[1] This section is based on the document: “Requirements and principles for a Future Media and 3D Internet,” created by the "Future Media and 3D Internet Task Force". Coordinated and supported by the Networked Media Unit of the DG Information Society and Media of the European Commission ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/netmedia/20090220-fid-rp-3-dg_en.pdf

creation of new applications along with the ability for multiple entities, which are implemented according to certain common rules, to communicate with each other (interoperability).

4 High-level FMI Network Architecture

We envision an FMI architecture which will consist of different virtual hierarchies of nodes (overlays), with different functionalities. In Figure 3, 3 layers are depicted, however this model would be easily scaled to multiple levels of hierarchy (even mesh instantiations, where nodes may belong to more than one layer) and multiple variations, based on the content and the service delivery requirements and constraints.

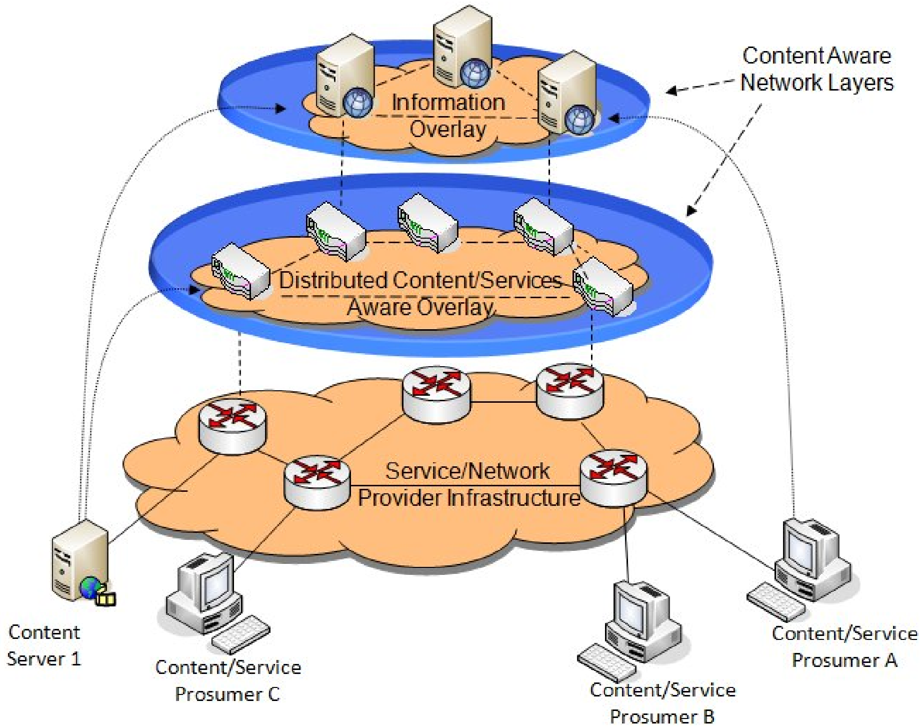

Figure 2: FMI high level architecture

In a realistic roll-out scenario, the FMI deployment is expected to be incremental. This is because we expect that today's existing legacy network nodes (core routers, switches, access points) will not only remain and will even be the majority for a number of years; thus the proposed architecture should be backwards compatible with current Internet deployment.

As shown in Figure 2, the ***Service/Network Provider Infrastructure Overlay*** is located at the lower layer. Users are considered as Content Producers (user generated content) and Consumers (we can then call them "Prosumers"). This Network Infrastructure Overlay is the service, ISP and network provider network infrastructure, which consists of nodes with limited functionality and intelligence (due to the cost of the network constraints). Content will be routed, assuming basic quality requirements and if possible and needed cached in this layer.

The medium layer is the ***Distributed Content/Services Aware Overlay.*** Content-Aware Network Nodes (e.g. edge routers, home gateways, terminal devices) will be located at this overlay. These nodes will have the intelligence to filter content and Web services that flow through them (e.g. via deep packet Inspection or signalling processing), identify streaming sessions and traffic (via signalling analysis) and provide qualification of the content. This information will be reported to the higher layer of hierarchy (***Information Overlay***). Virtual overlays (not shown in the figure) may be considered or dynamically constructed at this layer. We may consider overlays for specific purposes e.g. content caching, content classification (and depending on the future capabilities, indexing), network monitoring, content adaptation, optimal delivery/streaming.

With respect to content delivery, nodes at this layer may operate as hybrid client-server and/or peer-to-peer (P2P) networks, according to the delivery requirements. As the nodes will have information about the content and the content type/context that they deliver, hybrid topologies may be constructed, customised for streaming complex media such as Scalable Video Coding (SVC), Multi-view Video Coding (MVC).

At the highest layer, the Content/Services ***Information Overlay*** can be found. 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 P2P topology to secure corporate routers or even Data Centres in a distributed carrier-grade cloud network. The content may be stored/cached at the *Information Overlay* or at lower hierarchy layers. Though the *Information Overlay* we can be always aware of the content/services location/caching and the network information. Based on this information, a decision on the way that content will be optimally retrieved and delivered to the subscribers or inquiring users or services can be made.

4.1 Content production workflow

In the above layered architecture, the following steps for producing, publishing, caching, finding and consuming content may be considered.

The content is **produced** in two steps. It is **generated** and may be stored in a server, which can be for example a home server, a residential gateway, a network video server (e.g. YouTube) or even a mobile phone. In parallel, some metadata may be created manually by the content creator or generated automatically using annotation tools (or a combination of both). In all cases, the content is **published** inside the FMI architecture network either via a manual publishing procedure or via an automatic content discovery and identification procedure (via normal crawling or discovery "on-the-fly" using Deep Packet Inspection (DPI) techniques). The publish procedure enables the content to be found through the search engine and accessible through the FMI network[2].

A user can directly **consume** content utilizing the FMI network if he/she knows exactly the content item he/she is looking for. In most cases, users are looking for content items related to some abstract concepts or associated with some context; for example, a user can try to find videos of the goals of the Italian football team in the semi-final of the FIFA 2006 world cup. In order to consume content, the first step is to **search** for the content items related to this abstract concept using a search engine. The search engine will return content items, some of which may have been stored inside the FMI network. If the user desires to access content, stored/cached in the FMI overlay or not, it will **fetch** the content onto its local server. As the content is delivered in the FMI overlay network, it may also be **cached** by different cache nodes. This decision can be based on

[2] It should also be underlined that due to privacy issues and EC directives enforcements, the proposed FMI architecture may deal only with content that the creator/publisher has explicitly given the permission.

criteria such as the awareness of the content flowing through its nodes. Then, the user will receive and **consume** the content. As we assume that FMI will support high volumes of new media types which are more and more demanded by the users such as 3D videos, it is needed to provide content adaptation and enrichment based on the device used by the user, its connectivity, and any form of user interaction to cope with new and current media types distribution, assuring a compatibility between networks, terminals and contents irrespective of the media type.

The network will also **analyze** the content being requested through it. This information is collected and used by the search engine. In particular, the analysis determines the popularity of the content items and helps the content discovery (that would not be discovered by active crawling, or that would be discovered after the process end). It will also **monitor** its own state, as for example the connectivity between nodes, in order to report it to various components.

In the following, we will overview how these steps are implemented by the main architectural components.

5 FMI protocol stack

It should be noticed that the proposed protocol stack as described in this section, covers the complete FMI functionality which may exist at some edge nodes, servers and terminals but it is not expected to be present at all nodes of the network. On the contrary, due to cost limitations and the need for reusability of the existing infrastructure, grouping and subsets of the proposed functionality will be necessary based on the network planning.

For the FMI protocol stack, we propose the replacement of today's OSI protocol stack layers with functional stratums. We consider stratums as large vertical blocks, which target specific functionality and cover different viewpoints of the FI. As can be seen in Figure 3, for the FMIA we envision 4 stratums. At the lower part of the protocol stack, the ***network stratum*** is located, which assures to the higher layers that groups of bits (e.g. data packets, video chunks, control messages, out of band signalling) are delivered. This stratum is responsible for handling the network infrastructure related issues, including the large diversity and scale of nodes (i.e. routers, edge devices, residential gateways, servers, terminals and sensors), the virtual representations of network resources, the network topologies and load balancing and the physical characteristics of the network infrastructure.

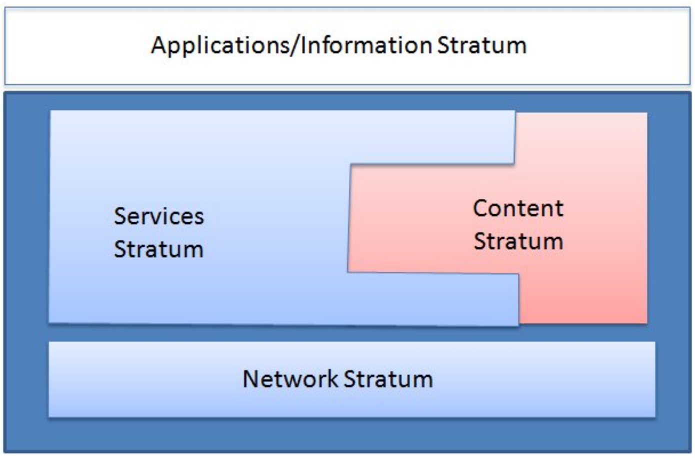

Figure 3: FMI proposed protocol stack

The FMI targets content creation, generation, delivery and consumption in the Future Internet environment. As such, we consider two stratums over the network stratum: the ***content stratum***, which assures efficient, timely and securely delivery of the content and offers a virtual representation of the content and the ***services stratum***, which assures flexible and reliable discovery and delivery of the services and the services' components that manipulate the content and the information. One should note that in the figure the services stratum embraces the content stratum. This is presented in this way in order to emphasise that in most cases the services components and APIs will be needed for content handling, though direct interfaces between the content and the network stratums is also considered. The network, service and content stratums are considered as the core of the FMI, while over them, we assume that the ***Applications and Information Stratum*** is located, which utilises the services and content stratums and interfaces with the end-user. Applications and information stratum is not considered as an integral part of FMIA, but the FMIA empowers and facilitates it.

5.1 Protocol Stack Functionalities

The FMI architecture protocol stack of Figure 3 is further analysed in Figure 4. We foresee that the Network Stratum will provide two basic functionalities: ***transmission*** of groups of bits and ***discovery*** of network resources (new nodes, nodes' capabilities, alarms, events, traffic analysis etc.). Moreover we consider that the transmission component will include issues like buffering and admission control, whereas the discovery component will include a sub-block offering semantics' analysis of the network resources. The network resources information is forwarded to the ***self-configuration*** & ***healing*** component, which offers sustainability, adaptability, continues self-organization and self-recovery and covers issues like mobility and network fluctuation. It is also responsible for monitoring and tracking activities (for accounting, traceability and security issues) and realising policies of the network operator, supporting tussle and open competition. The interface between the network stratum and the services/content components is the ***abstraction sub-layer***, which offers a unified representation of the network resources to the top stratums. This approach has been chosen to hide the complexity of the low-level stratum resources to the higher-level ones, and ease the operation of the applications and information exchange.

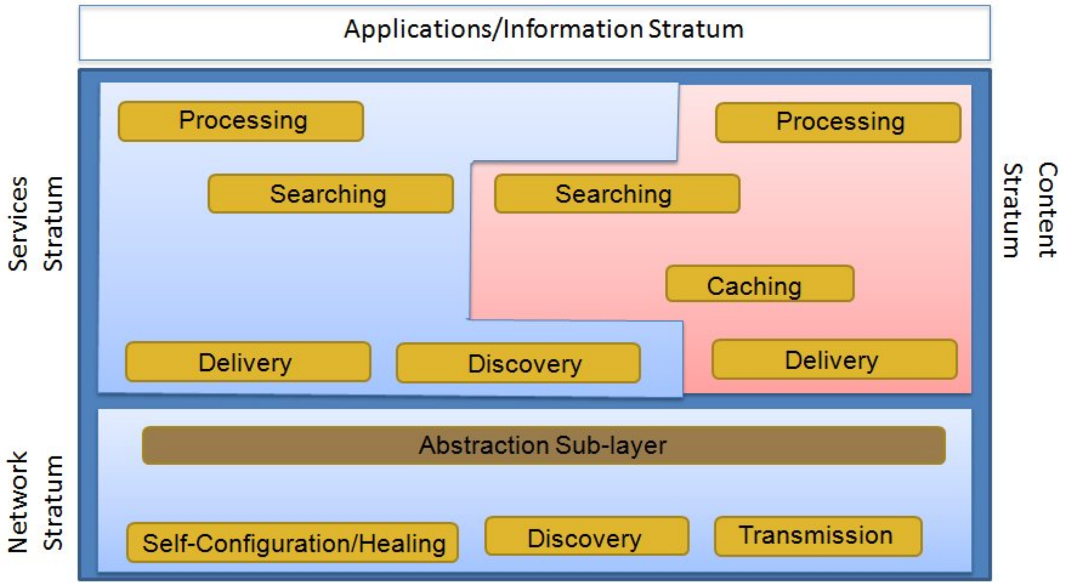

Figure 4: Detailed FMI proposed protocol stack

In the services stratum, we foresee ***service resources searching*** (active) and ***discovery*** (passive) of any kind of service, service component or even process that offers an open interface (e.g. Semantic Web services, RESTful services, WSDL services, processes, services mash-ups, etc.), including semantics analysis of the discovered service components. The services stratum also offers a ***services processing*** component, which handles services ranking, reasoning, dynamic composition and execution. The processing component is also responsible for authentication and authorization, monitoring and tracking activities (for accounting, traceability and security issues) and realising policies of the service providers, supporting tussle and open competition. It is supported by the discovery and the services ***searching*** components. Finally, services' related communication is achieved in a reliable, transparent, efficient and secure manner via the ***services delivery*** component. This may be in the simplest case extensions to IP routing or any other routing protocol specialised for service delivery.

Finally in the content stratum, one of the most important foreseen units is the ***content delivery*** component. The content delivery component is responsible for specialised, reliable, transparent, efficient, timely and secure content-aware delivery. Many design orientations and solutions are candidate for this component, including IP routing, routing by name [11], P2P delivery, adaptive streaming, content hybrid routing [12], publish-subscribe mechanisms, etc. but we have preferred to keep it open for future developments, assuring the compatibility and openness. Moreover, the delivery component in some cases may penetrate the virtualization sub-layer in order to directly communicate with the transmission component. Content delivery will operate in very close collaboration with the ***content caching*** component, which will provide temporal storage of the delivered content, both as autonomous content files and as part of video streams (content chunks). The content caching components from different nodes may collaborate in order to provide a fully distributed storage space, following various distributed caching and cloud technologies [13]. The ***content processing*** component will offer coding, encryption, adaptation and enrichment functions in the network. The processing, caching and delivery components will be assisted by the ***searching*** component, which will offer content discovery, identification and content popularity calculation, both assisting towards caching optimization and for monitoring and tracking activities.

6 FMI network architecture

As already explained, due to network planning cost limitations and the need for reusability of the existing infrastructure, it is expected that different nodes in the network may not host all stratums and/or host subsets of the proposed functionality of each stratum. Based on this assumption, Figure 5 shows a hierarchical view of the FMI network architecture, which details the network architecture of Figure 4. The main functionality of FMI resides in the content and services distributed overlay, where we have defined the following functional modules/entities:

- **Delivery Nodes**: They are responsible for the content & services delivery, IP acceleration and efficient content streaming (including P2P overlays creation).
- **Caching Nodes**: They are responsible for content caching, caching optimization and content replacement in collaboration with the cache content optimization entity.

- **Discovery Nodes**: They contribute to the discovery of new and calculating the popularity of known services and content (stored or streaming). They also measure traffic analytics and help towards network topology discovery.
- **Process Nodes:** They are responsible for services processing in-the network and content adaptation & enrichment.

An assumption would be that delivery and caching nodes' functionality would co-exist in most cases, followed by the discovery and the processing functionality.

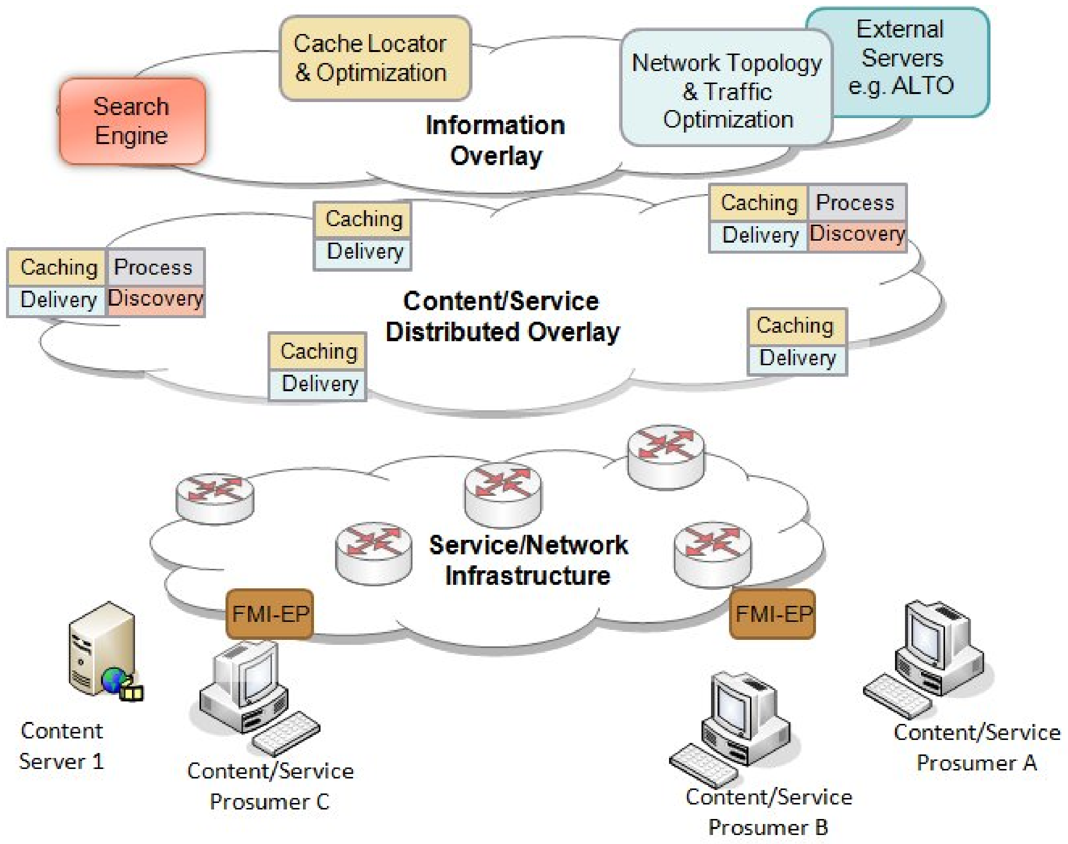

Figure 5: FMI network architecture

The proposed FMI functionality may be fully distributed at the content/services distributed overlay. For our explanation, for presentation and simplicity reasons, we may assume that some functionality is provided by an additional ***Information Overlay***, which handles the following functional modules/entities:

- **Search Engine**: It is a distributed system that discovers and indexes the content and the services, processes the queries from the users and returns relevant results ordered according to several criteria. It may also be considered at an application overlay.
- **Content Cache Locator & Optimizer:** This entity may exist as a group of dedicated physical nodes or may be a fully distributed abstract functionality. The locator module will redirect content requests to the "best" cached copy, where "best" is defined based on perceived Quality of Service (PQoS) of the user. In order to make the decision it may also communicate with the network/traffic monitor entity. The optimizer module will support caches in deciding which object they should store or evict.

- **Network Topology/Traffic Optimizer**: It is responsible for gathering all network related information: topology, traffic, characteristics of the user Internet access and optionally user location. It may be a variation of an IETF ALTO server [14] or communicate with/supported by external traffic and network optimizer servers.

Finally, as entry points to the FMI we have defined the ***FMI Entry Point (FMI-EP)***. The FMI-EP may be hosted at a local router or a Residential Gateway and is responsible for seamless operation, termination of FMI protocol stack processes (e.g. receiving and adapting content delivery) and optimal content fetching and streaming.

One may notice that some functionality could be aggregated in less functional entities or that some entities could be removed. For example, the FMI-EP module may be overloaded to perform also the Content Cache Locator role, whereas the Cache optimizer would be distributed at the overlay network. Indeed, this may depend on the final implementation approach chosen (as the purpose of this section was to emphasize the functional blocks, rather than propose an actual instantiation).

7 Conclusions

The immense success of Internet has created high hopes and expectations for new immersive and real-time applications and services. Towards this future environment there are supporters that believe that re-engineering of some Internet protocols is adequate and others that believe that a holistic re-design of the Internet is needed.

We believe that extensions, enhancements and re-engineering of today’s Internet protocols may solve many of problems which have appeared in the last decade. But the Future Internet design is a multi-dimensional problem. Improvements in each dimension may help to create new solutions and solve some of the partial problems, but a holistic approach is needed to provide a holistic solution which can really cover the expected applications and services described in this paper.

One of the possible solutions is the reference architecture, its functions and the new protocol stack that targets the Future Media Internet as a holistic ecosystem as we have described, with the advantages expressed in this paper, and the open opportunities to accommodate different instantiations which may appear in the future.

8 References

[1] http://www.nsf.gov/pubs/2010/nsf10528/nsf10528.htm

[2] www.geni.net/netse_about.html

[3] www.nets_find.net

[4] www.geni.net/netse_about.htm

[5] akari-project.nict.go.jp/eng/overview.htm

[6] mmlab.snu.ac.kr/fiw2007/presentations/architecture_tschoi.pdf

[7] http://www.future-internet.eu/

[8] http://www.fi-nextmedia.eu

[9] Th. Zahariadis, F. Junqueira, L. Celetto, E. Quacchio, S. Niccolini, P. Plaza, “Content aware searching, caching and streaming,” 2nd International Conference

on Telecommunications & Multimedia, Chania, Greece, 14-16 July 2010, pp. 263-270

[10] M. Alduán, F. Álvarez, Th. Zahariadis, N. Nikolakis, F. Chatzipapadopoulos, D. Jiménez, J. Manuel Menéndez, “Architectures for Future Media Internet,” 2nd International Conference on User Centric Media, Palma de Mallorca, Sept. 1-3, 2010

[11] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, R. Braynard, “Networking Named Content,” Proceeding of ACM CoNEXT 2009. Rome, Italy, December 2009

[12] M. Serafini, et.al., “D2.2: End-to-End Future Content Network Specification,” COAST Consortium, August 2010

[13] The CORAL Content Distribution Network,” http://www.coralcdn.org/overview/

[14] https://datatracker.ietf.org/wg/alto/charter/

[15] Clark, D.D.; Wroclawski, J.; Sollins, K.R.; Braden, R., “Tussle in cyberspace: defining tomorrow's Internet”, IEEE/ACM Transactions on Networking, Vol 13, Issue 3, pages 462-475, June 2005.

9 FMIA-TT members

Name	Affiliation	Contributor
Jan Bouwen	Alcatel-Lucent	X
Pablo Cesar	CWI: Centrum Wiskunde & Informatica	X
Van Jacobson	Palo Alto Research Centre (PARC)	
George Pavlou	University College London (UCL)	X
Pablo Rodriguez	Telefonica R&D	
Nikos Laoutaris	Telefonica R&D	X
Gerardo Garcia	Telefonica R&D	
Giovanni Pau	University of California, Los Angeles	X
Christian Timmerer	Klagenfurt University	X
Olivier Festor	INRIA	X
Gonzalo Camarillo	Ericsson Research	X
Marcelo Bagnulo	Carlos III University	
Thomas Steiner	Google	X
Amar-Djalil Mezaour	Exalead	X
Sergios Soursos	Intracom Telecom	X
Theodore Zahariadis (Editor)	Synelixis	X
Thanasis Tsiodras	Synelixis	X
Petros Daras	CERTH	X
Paul Moore	Atos Origin	X
Ebroul Izquierdo	Queen Mary University London	X
Tomas Piatrik	Queen Mary University London	X
Federico Alvarez	Technical University of Madrid	X
Maria Alduan	Technical University of Madrid	X
Isidro Laso	European Commission	

Table 1. Future Media Internet Architecture Think Tank - Members list