

Internet Engineering Task Force
Internet Draft
Intended Status: Informational
Expires: April 26, 2012

Th. Zahariadis, Ed.
Synelixix
Y. Le Louedec
Orange Labs
Ch. Timmerer
UNI-KLU
S. Spirou
Intracom
D. Griffin
UCL
October 24, 2011

**Catalogue of Advanced Use Cases for
Content Delivery Network Interconnection**

[draft-fmn-cdni-advanced-use-cases-00](#)

Abstract

The purpose of this draft is to complement the current CDNi WG use-cases with a catalogue of advanced, longer-term CDN interconnection use cases. The work has been the contribution of six European Commission (EC) co-funded projects, which are part of the EC Future Media networks (FMN) cluster.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process.

Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1	Introduction	4
1.1	Terminology	4
1.2	Abbreviations	4
2	Advanced/Long-term CDNI Use Cases	5
2.1	Use Case 1: Caching-CDN interconnection	5
2.2	Use Case 2: CDN-CDN interconnections at large scale	6
2.3	Use Case 3: Dynamic adaptive streaming over HTTP in multi-CDNs	7
2.4	Use Case 4: Dynamic expansion of CDN capacity and geographical reach	8
3	Relationship between CDNI and Information-Centric Networking	9
4	Acknowledgements	11
4.1	List of Contributors	12
5	References	12
5.1	Normative References	12
5.2	Informative reference	13
	Authors' Addresses	14

1 Introduction

The purpose of this draft is to complement the current CDNi Work Group short-term use-cases with a catalogue of advanced, longer-term CDN interconnection use cases. The proposed catalogue of use cases is coming from or inspired by work in the European Commission (EC) Future Media Network (FMN) cluster research projects. Though they are beyond the current short-term objectives of the CDNi WG, the proposed use cases could drive future work in the CDNi WG, especially in case of WG re-chartering (once the present goals will be reached). The use cases are derived from ongoing research work in six EC co-funded projects where concrete design, implementation and evaluation work is being undertaken to validate the approaches.

Moreover, this draft compares CDNs with Information-Centric Networking (ICN) which have similar goals and use cases. We discuss here this relation for the benefit of people and organizations working in these areas.

1.1 Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119](#) [[RFC2119](#)].

The present document reuses terminology defined by CDNi WG documents [[I-D.ietf-cdni-problem-statement](#)], [[I-D.ietf-cdni-use-cases](#)] and [[I-D.ietf-cdni-requirements](#)].

1.2 Abbreviations

[Ed. Note: List of abbreviations to be updated later]

- o CSP: Content Service Provider
- o dCDN: downstream CDN
- o FMN: Future Media Networks
- o ICN: Information-Centric Networking
- o NSP: Network Service Provider
- o QoE: Quality of Experience
- o SLA: Service Level Agreement
- o uCDN: upstream CDN

- o MPD: Media Presentation Description
- o DASH: Dynamic Adaptive Streaming over HTTP

2 Advanced/Long-term CDNi Use Cases

This draft includes four advanced CDNi use cases. The first use case introduces a Caching-CDN interconnection that put emphasis on the in-the network caching mechanism. As CDNi WG mainly targets interconnection between two CDNs, the second use case extends this use case to large-scale CDNs. The third use case introduces support of dynamic adaptive streaming over HTTP (DASH) in a multi-CDNs context, which may be today's most promising video distribution method as it overcomes limitations posed by firewalls and promises efficient distribution utilizing CDNs and network caching resources. Finally, the forth use case proposes the dynamic expansion of CDN capacity and geographical reach.

2.1 Use Case 1: Caching-CDN interconnection

Some telecom operators and NSP deploy caching servers, a.k.a. "transparent caching" servers, in their IP networks in order to cache content by CDNs over Internet, with as goal to relax and balance the load on their IP networks.

Today in most situations the two overlay networks - the upstream CDN (uCDN) and the downstream CDN (dCDN) set of transparent caching servers - run independently from each other. There is no specific interconnection interface between the two overlay networks.

There could be some interest to set up interfaces between these overlay networks (of course on condition that their owners get the right business incentives for).

Here are two illustrations of the potential benefits (this is not an exhaustive list):

- Setting up a "logging interface" between the two overlay networks could be beneficial for providing the uCDN (and beyond the Content Service Providers, CSP) with useful logging, monitoring, reporting data.
- Setting up a "CDN metadata interface" between the two overlay networks could be beneficial for allowing the uCDN to request that a given content file be purged from, or invalidated in, any downstream caching server.

The current charter of the CDNi WG states "the WG will not define

support for transparent caching across CDNs". The first priority is indeed to meet the goals and milestones specified in the current charter.

This use case proposal aims at recommending that this "caching-cdn interconnection" use case be considered in case of WG re-chartering (once the present goals will be reached).

The first difference with the present cdn-cdn interconnection model addressed by the CDNi WG is that there is no request routing interface in this "caching-cdn interconnection" use case. Then different sub-cases can be envisioned, depending on which CDNi interfaces are leveraged (with or without logging interface, with or without CDNI metadata interface, etc.). In a first approach, this "caching-cdn interconnection" use case could therefore be considered as a sub-case of the current CDNi WG's cdn-cdn interconnection model.

Note. Once such specific interfaces are set up between the uCDN and the dCDN set of transparent caching servers, the latter ones cannot be any more considered as fully transparent, at least from the viewpoint of the uCDN. This should call for a slight evolution of the terminology.

2.2 Use Case 2: CDN-CDN interconnections at large scale

The current focus of the CDNi WG is on the interconnection between two CDNs.

The purpose here would be to investigate situations where (possibly much) more than two CDNs are involved in a CDN federation.

As stated by the Amplification Principle defined in [[RFC3439](#)] "there do exist non-linearities which do not occur at small to medium scale, but occur at large scale". As a result, the number of involved CDNs could impact on the requirements and constraints, especially in terms of scalability, and therefore on the conceivable technical options to ensure interconnection.

As an illustration of the amplification principle application in CDNi, the initial considerations, and potential issues, about request routing in CDN interconnect scenarios presented in [I-D.stiemerling-cdni-routing-cons] could be even more critical in large scale CDN federations.

This would possibly lead to the definition of specific sub cases corresponding to cdn-cdn interconnections at large scale. Yet the

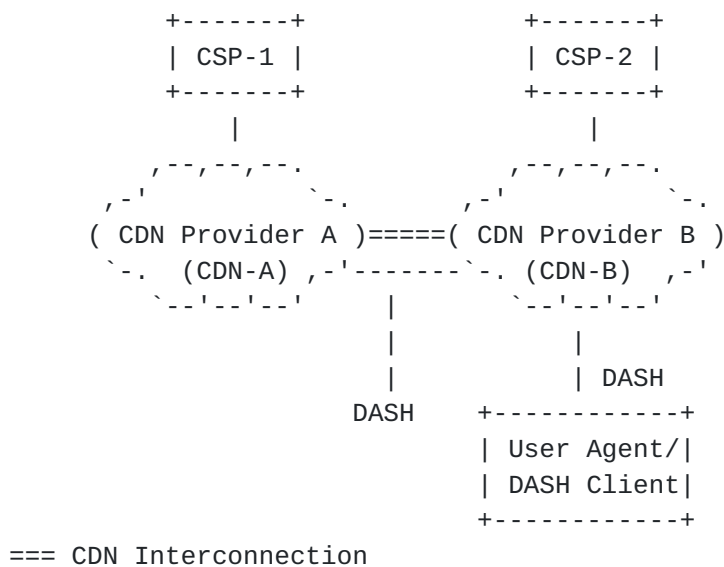
first priority (and prerequisite before investigating such large scale use cases) is to meet the goals and milestones specified in the current charter, i.e. to specify adequate interfaces for interconnecting two CDNs. So this proposal aims at recommending that this "cdn-cdn interconnections at large scale" use case be considered in case of WG re-chartering (once the present goals will be reached).

2.3 Use Case 3: Dynamic adaptive streaming over HTTP in multi-CDNs

The ever-growing video traffic on the Internet requires more efficient content distribution mechanisms. Compared to existing RTP/RTSP-based streaming or HTTP progressive download, Dynamic Adaptive Streaming over HTTP (DASH) enables stateless communication between a client and the corresponding server, better content adaptability and support for live media services [[Stockhammer2011](#)]. Intrinsic characteristic of DASH is the content fragmentation into various representations comprising segments (or sub-segments) of different encodings of one or several media components. These components/segments are transferred, along with content metadata descriptors referred to as media presentation description (MPD), to the origin servers. Using MPD, clients request the segments utilizing HTTP GET or partial GET requests.

DASH is considered as a promising solution for efficient and high-quality delivery of streaming services over the Internet and is currently standardized as 3GP-DASH, MPEG-DASH, and OIPF HAS. It supports among others, re-use of existing technologies (codecs, containers, etc.), deployment on top of HTTP-CDNs, re-use of HTTP origin and cache/proxy servers, re-use of existing media play-out engines, as well as on-demand, live and time-shifted delivery.

For example, DASH may be exploited within specific CDNs enabling clients to retrieve content hosted by given surrogates, taking into account the scale, the coverage and the reliability of HTTP-based CDN systems. Additionally, DASH can be also utilized as a delivery solution in a CDN Interconnect environment, enabling content acquisition among different providers. In such a case, a content service provider (e.g., CSP-1 in the following figure) will first ingest the prepared content into CDN-A, so that each surrogate can act as a DASH server offering the streaming service to the requesting End-User, acting as DASH client, connected with a different CDN (e.g., CDN-B). Towards this, CDN Interconnection requires interfaces among CDNs to be capable of facilitating their collaboration besides ensuring content streaming efficiently.



In other words, in an interconnected CDN environment the definition of a control interface is required, which will enable DASH clients to acquire specific information from a hosting CDN over multiple-CDNs. Towards these, the anticipated interface must be able of providing/delivering:

- The Media Presentation Description that contains metadata to construct appropriate HTTP-URLs to access segments and to provide the streaming service to the user.
- The Number of source surrogates: one or multiple (content segments can be acquired from multiple sources).
- The location of source surrogates: e.g., locality-aware capabilities for choosing the nearest source(s), as a matter of geographical (internal or external) or network-based metrics (e.g., using latency or cost-based metrics).
- Delivery protocols: Definition of the delivery protocols used for the delivery of segments.
- Additional metadata for content delivery (e.g., metadata for content-aware networks).

2.4 Use Case 4: Dynamic expansion of CDN capacity and geographical reach

ISPs may offer a set of specialized network services which may be invoked by their customers, including CDN providers, with appropriate prior agreements and possible payments. The network service of most relevance to this use case is the provision of caching resources located within the ISP. A CDN provider making use of such a service

may invoke new caching resources within a local or remote ISP to dynamically create new CDN surrogate nodes. The newly created resources are provided by the network provider but under the control of the CDN provider.

This is an alternative way of dealing with increased load on a CDN, and a CDN provider who is able to invoke dynamic nodes is therefore able to expand dynamically to accommodate increased traffic demand or extend geographical reach. Such a CDN provider could a) advertise its elasticity to upstream CDNs which could simplify/enhance its resource-routing algorithm decisions, or b) advertise its new capabilities dynamically as it expands/contracts. These announcements could be made part of the CDNI control interface interactions and contributions could be made on the protocol extensions to announce capabilities dynamically and also to propose elasticity metrics.

3 Relationship between CDNI and Information-Centric Networking

CDN interconnection has similar goals and use cases to Information-Centric Networking (ICN). We discuss here this relation for the benefit of people and organizations working in these areas. We start with a very brief description of CDNs and ICN in order to have a common understanding. We then discuss similarities and differences in terms of objectives, technical approach, deployment and business models. Finally, we explore the possibility of interaction and coexistence of ICN and CDN in a future Internet. CDNs are real working systems, while ICN is still at the research stage. It is important to note that our analysis is based on the state of ICN research and the assumption that ICN will meet its design goals.

A CDN is a privately owned overlay network that aims to optimize delivery of content from (Content Service Providers) CSPs to End Users. CDNs are independent of each other. Optimization is in terms of performance, availability and cost. CDNs operate by serving content to User Agents from managed caches - called surrogates - at the edge of the network. CDNs may also use reserved network resources. The CDN provider collaborates with NSPs on surrogate placement and network resource reservation. The deployment, extension and (part of) the operation of CDNs are centrally managed. To maintain transparency for End Users, normal content locators (e.g., URLs) are used to access content. However, the CDN treats those locators as simple content identifiers when selecting a surrogate, in effect, decoupling the locator from the content location.

ICN shares many of the goals of CDNs. Optimization of delivery with ICN usually entails caching, QoS-aware routing and traffic differentiation, to various degrees. Unlike CDNs, ICN aims at

operating natively at the NSP level (although operation as an overlay is also possible). An NSP deploys ICN-enabled elements (mainly routers) with minimal planning in terms of content demand. The ICN reach grows with each new ICN-enabled NSP, without any central management. For content access, ICN uses explicit content identifiers that are independent from the content location.

A main objective of both CDNs and ICN is the effective delivery of content from CSPs to End Users. ICN also aims at accommodating End Users as small CSPs, i.e., End Users who want to share content. The decoupling of content identifiers from the content location is an explicit goal in ICN, while in CDN this is a by-product.

The main elements of a CDN are a set of content servers and a request redirection system. The content servers are carefully placed to be close to the User Agents and can be populated prior to any requests based on foreseen demand. An End User requests content with a normal URL, which triggers a part of the redirection system that resides at the seeming location of the content. The redirection system selects a content server based on criteria aiming to optimize some aspect of the delivery task. The selected content server then delivers the requested content to the User Agent.

In ICN, any edge router with storage capabilities can act as a content server, which is populated based on actual demand. Alternatively, or in addition, QoS-aware routing and traffic differentiation techniques are used to establish a path from the content source to the User Agent. Content is requested with a dedicated identifier. This identifier can be used to route the request to the content source or to an intermediate content server, while constructing the path. Alternatively, the identifier can be used as an index in a directory system to retrieve content metadata. The metadata are then used to setup a path from the content source to the User Agent.

A CDN constitutes an administrative domain, whereas in ICN an administrative domain can be as granular as an ICN router. When viewed in terms of administrative domains, CDN interconnection resembles the interconnection of ICN elements/domains. As such, request routing in CDNI and in ICN presents similar technical challenges. Following this thinking, the Content Distribution Metadata and CDNI Metadata of CDNI become related to the content metadata that are used in ICN to establish a path.

A CDN is deployed as an overlay with all its elements independent from NSPs and belonging to the same CDN Provider. The CDN Provider works closely with the CSP in order to ingest and place content. There's also close collaboration between the CDN Provider and the

NSPs for server placement and reservation of resources. Once the CDN is operational, any unforeseen change in CSP requirements, network conditions or End User demand will force the CDN Provider to re-design the deployment. To avoid this, CDNs are designed with over-provisioning and redundancy. On the other hand, ICN is deployed as NSP infrastructure. Each NSP owns and independently manages the ICN elements in its domain. There's no "ICN Provider" to oversee the deployment of ICN. The system grows as each NSP becomes ICN-enabled. In order to use ICN, CSPs need to provide content source servers and content metadata. ICN is designed with flexibility in mind, which permits coping with many unforeseen events.

In a typical business case for CDNs, a CDN is deployed in an NSP domain and the NSP acts both as a CDN Provider and a (sole) CSP. NSPs use this model to offer content services, such as IPTV, to their End Users and so to differentiate their business. In another model for CDNs, a CDN Provider agrees with several NSPs to deploy content servers and reserve resources in their domain. The CDN Provider then sells content delivery services to interested CSPs.

The ICN business model is very similar to that for connectivity services. An NSP deploys ICN in its domain and makes transit or peering agreements with other ICN-enabled NSPs. The NSP then offers content delivery services to CSPs. It is also possible for the NSP to offer content consumption services to its End Users.

In a future Internet where ICN is deployed by some NSPs, it would be difficult to meet end-to-end the ICN performance goals, because of the "holes" between ICN domains. In such a case, CDNs can act as overlay bridges, because they might already have content close to the downstream ICN domain. In essence, from the ICN point of view, CDN Providers would become CSPs. However, the incentive for the CDN Provider is unclear, as such a move could undermine its own content delivery service. Borrowing from the motivation for CDNI, a CDN Provider who wants to extend its reach, instead of interfacing with a neighboring CDN, could interface with an ICN-enabled NSP. This is of course a scenario that only market forces and time can validate.

4 Acknowledgements

This draft has been produced by the Future Media Networks (FMN) cluster of the Networked Media Systems FP7 projects. The work leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) in (alphabetic order):
ALICANTE (FP7-ICT-248652; <http://www.ict-alicante.eu/>),
COAST (FP7-ICT-248036; <http://www.coast-fp7.eu/>),
COMET (FP7-ICT-248784; <http://www.comet-project.org/>),
ENVISION (FP7-ICT-248565; <http://www.envision-project.org/>),

NEXTMEDIA (FP7-ICT-249065; <http://www.fi-nextmedia.eu/>) and
OCEAN (FP7-ICT-248775; <http://www.ict-ocean.eu/>).

4.1 List of Contributors

Andrzej Beben, Warsaw University of Technology, Poland;
Christian Timmerer, UNI-KLU, Austria;
David Florez Rodriguez, Telefonica R&D, Spain;
David Griffin, Yiannis Psaras, Raul Landa, UCL, UK;
Daniel Negru, Labri, France;
Evangelos Pallis, Petros Anapliotis, TEIC, Greece;
George Xilouris, DEMOKRITOS, Greece;
Isidro Laso, European Commission, Belgium (on a personal basis);
Klaus Satzke, Alcatel-Lucent Bell Labs, Germany;
Michalis Georgiadis, PrimeTel, Cyprus;
Ning Wang, University of Surrey, UK;
Spiros Spirou, Intracom Telecom, Greece;
Theodore Zahariadis (editor), Synelixis, Greece;
Yannick Le Louedec, Orange Labs, France;
Yiping Chen, Daniel Negru, CNRS-LaBRI, France.

5 References

5.1 Normative References

- [RFC3439] R. Bush, D. Meyer, "Internet Architectural Guidelines,"
[RFC 3439](#), <http://www.ietf.org/rfc/rfc3439.txt> (updates
[RFC 1958](#)), December 2002.
- [I-D.ietf-cdni-problem-statement] Niven-Jenkins, B., Faucheur, F.,

and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", [draft-ietf-cdni-problem-statement-00](#) (work progress), September 2011.

[I-D.ietf-cdni-requirements] Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", [draft-ietf-cdni-requirements-00](#) (work in progress), September 2011.

[I-D.ietf-cdni-use-cases] Bertrand, G., Stephan, E., Watson, G., Burbridge, T., Ma, K., "Use Cases for Content Delivery Network Interconnection", [draft-ietf-cdni-use-cases-00](#) (work in progress), September 2011.

[5.2](#) Informative reference

[I-D.stiemerling-cdni-routing-cons] Stiemerling, M., "Considerations on Request Routing for CDNI", [draft-stiemerling-cdni-routing-cons-00](#), July 2011

[Stockhammer2011] T. Stockhammer, "Dynamic adaptive streaming over HTTP: standards and design principles", ACM Proceedings of the second annual ACM conference on Multimedia systems, 2011.

[3GP-DASH] ETSI TS 126 247 v10.0.0 (2011-06) Universal Mobile Telecommunications System (UMTS); LTE; Transparent end-to-end Packet-switched Streaming Service (PSS); Progressive Download and Dynamic Adaptive Streaming over HTTP (3GP-DASH) (3GPP TS 26.247 version 10.0.0 Release 10).

Authors' Addresses

Yannick Le Louedec
Orange Labs, France
Email: yannick.loulouedec@orange.com

Christian Timmerer
UNI-KLU, Austria
Email: christian.timmerer@itec.uni-klu.ac.at

Spiros Spirou
Intracom, Greece
Email: spis@intracom.com

David Griffin
UCL, UK
Email: dgriffin@ee.ucl.ac.uk

Theodore Zahariadis
Synelixis, Greece
Email: zahariad@synelixis.com

