

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 25, 2012

B. Niven-Jenkins
Velocix (Alcatel-Lucent)
F. Le Faucheur
Cisco
N. Bitar
Verizon
June 23, 2012

Content Distribution Network Interconnection (CDNI) Problem Statement draft-ietf-cdni-problem-statement-07

Abstract

Content Delivery Networks (CDNs) provide numerous benefits: reduced delivery cost for cacheable content, improved quality of experience for End Users and increased robustness of delivery. For these reasons they are frequently used for large-scale content delivery. As a result, existing CDN Providers are scaling up their infrastructure and many Network Service Providers (NSPs) are deploying their own CDNs. It is generally desirable that a given content item can be delivered to an End User regardless of that End User's location or attachment network. This is the motivation for interconnecting standalone CDNs so they can interoperate as an open content delivery infrastructure for the end-to-end delivery of content from Content Service Providers (CSPs) to End Users. However, no standards or open specifications currently exist to facilitate such CDN interconnection.

The goal of this document is to outline the problem area of CDN interconnection for the IETF CDNI (CDN Interconnection) working group.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on December 25, 2012.

Copyright Notice

Copyright (c) 2012 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.

Table of Contents

| | | |
|-----------------------------|--|--------------------|
| 1. | Introduction | 5 |
| 1.1. | Terminology | 6 |
| 1.2. | CDN Background | 10 |
| 2. | CDN Interconnection Use Cases | 11 |
| 3. | CDN Interconnection Model & Problem Area for IETF | 12 |
| 4. | Scoping the CDNI Problem | 16 |
| 4.1. | CDNI Request Routing Interface | 17 |
| 4.2. | CDNI Metadata Interface | 17 |
| 4.3. | CDNI Logging Interface | 18 |
| 4.4. | CDNI Control Interface | 18 |
| 5. | IANA Considerations | 18 |
| 6. | Security Considerations | 18 |
| 6.1. | Security of the CDNI Control interface | 19 |
| 6.2. | Security of the CDNI Request Routing Interface | 19 |
| 6.3. | Security of the CDNI Metadata interface | 20 |
| 6.4. | Security of the CDNI Logging interface | 20 |
| 7. | Acknowledgements | 20 |
| 8. | References | 20 |
| 8.1. | Normative References | 20 |
| 8.2. | Informative References | 20 |
| Appendix A. | Design considerations for realizing the CDNI Interfaces | 23 |
| A.1. | CDNI Request Routing Interface | 23 |
| A.2. | CDNI Metadata Interface | 25 |
| A.3. | CDNI Logging Interface | 26 |
| A.4. | CDNI Control Interface | 27 |
| Appendix B. | Additional Material | 28 |
| B.1. | Non-Goals for IETF | 28 |
| B.2. | Relationship to relevant IETF Working Groups & IRTF Reserach Groups | 29 |
| B.2.1. | ALTO WG | 29 |
| B.2.2. | DECADE WG | 30 |
| B.2.3. | PPSP WG | 31 |
| B.2.4. | IRTF P2P Research Group | 32 |
| Appendix C. | Additional Material | 32 |
| C.1. | Related standardization activites | 32 |
| C.1.1. | IETF CDI Working Group (Concluded) | 33 |
| C.1.2. | 3GPP | 34 |
| C.1.3. | ISO MPEG | 34 |
| C.1.4. | ATIS IIF | 35 |
| C.1.5. | CableLabs | 35 |
| C.1.6. | ETSI MCD | 35 |
| C.1.7. | ETSI TISPAN | 36 |
| C.1.8. | ITU-T | 36 |
| C.1.9. | Open IPTV Forum (OIPF) | 36 |
| C.1.10. | TV-Anytime Forum | 37 |

| | | |
|-------------------------|--|--------------------|
| C.1.11. | SNIA | 37 |
| C.1.12. | Summary of existing standardization work | 37 |
| C.2. | Related Research Projects | 39 |
| C.2.1. | OCEAN | 39 |
| C.2.2. | Eurescom P1955 | 39 |
| Authors' | Addresses | 40 |

1. Introduction

The volume of video and multimedia content delivered over the Internet is rapidly increasing and expected to continue doing so in the future. In the face of this growth, Content Delivery Networks (CDNs) provide numerous benefits: reduced delivery cost for cacheable content, improved quality of experience for End Users (EUs) and increased robustness of delivery. For these reasons CDNs are frequently used for large-scale content delivery. As a result, existing CDN Providers are scaling up their infrastructure and many Network Service Providers (NSPs) are deploying their own CDNs.

It is generally desirable that a given content item can be delivered to an EU regardless of that EU's location or the network they are attached to. However, a given CDN in charge of delivering a given content may not have a footprint that expands close enough to the EU's current location or attachment network, or may not have the necessary resources, to realize the user experience and cost benefit that a more distributed CDN infrastructure would allow. This is the motivation for interconnecting standalone CDNs so that their collective CDN footprint and resources can be leveraged for the end-to-end delivery of content from Content Service Providers (CSPs) to EUs. As an example, a CSP could contract with an "authoritative" CDN Provider for the delivery of content and that authoritative CDN Provider could contract with one or more downstream CDN Provider(s) to distribute and deliver some or all of the content on behalf of the authoritative CDN Provider.

A typical end to end content delivery scenario would then involve the following business arrangements:

- o A business arrangement between the EU and his CSP, authorizing access by the EU to content items controlled by the CSP.
- o A business arrangement between the CSP and an "authoritative" CDN Provider where the CSP authorizes the CDN Provider to perform the content delivery on behalf of the CSP.
- o A business arrangement between the authoritative CDN Provider and another (or other) CDN(s) where the authoritative CDN may delegate the actual serving of some of the content delivery requests to the other CDN(s). A particular case, is where this other CDN Provider happens to also be the Network Service Provider providing network access to the EU, in which case there is also a separate and independent business relationship between the EU and the NSP for the corresponding network access.

The formation and details of any business relationships between a CSP and a CDN Provider as well as between one CDN Provider and another CDN Provider are out of scope of this document. However, this

document concerns itself with the fact that no standards or open specifications currently exist to facilitate such CDN interconnection from a technical perspective.

One possible flow for performing an end to end content delivery across a CDN Interconnect is described below:

- o The initial request from an EU's User Agent is first received by the authoritative (upstream) CDN, which is the CDN with a business arrangement with the CSP.
- o The authoritative (upstream) CDN may serve the request itself, or it may elect to use CDN Interconnect to redirect the request to a downstream CDN that is in a better position to do so (e.g. a downstream CDN that is "closer" to the EU).
- o The EU's User Agent will "follow" the redirect returned by the authoritative CDN and request the content from the downstream CDN. If required the downstream CDN will acquire the requested content from the authoritative (upstream) CDN, and if necessary the authoritative CDN will acquire the requested content from the Content Service Provider.

The goal of this document is to outline the problem area of CDN interconnection. [Section 2](#) discusses the use cases for CDN interconnection. [Section 3](#) presents the CDNI model and problem area being considered by the IETF. [Section 4](#) describes each CDNI interface individually and highlights example candidate protocols that could be considered for reuse or leveraging to implement the CDNI interfaces. [Appendix B.2](#) describes the relationships between the CDNI problem space and other relevant IETF Working Groups and IRTF Reserach Groups.

[1.1. Terminology](#)

This document uses the following terms:

Content: Any form of digital data. One important form of Content with additional constraints on distribution and delivery is continuous media (i.e. where there is a timing relationship between source and sink).

Metadata: Metadata in general is data about data.

Content Metadata: This is metadata about Content. Content Metadata comprises:

1. Metadata that is relevant to the distribution of the content (and therefore relevant to a CDN involved in the delivery of that content). We refer to this type of metadata as "Content

Distribution Metadata". See also the definition of Content Distribution Metadata.

2. Metadata that is associated with the actual Content or content representation, and not directly relevant to the distribution of that Content. For example, such metadata may include information pertaining to the Content's genre, cast, rating, etc as well as information pertaining to the Content representation's resolution, aspect ratio, etc.

Content Distribution Metadata: The subset of Content Metadata that is relevant to the distribution of the content. This is the metadata required by a CDN in order to enable and control content distribution and delivery by the CDN. In a CDN Interconnection environment, some of the Content Distribution Metadata may have an intra-CDN scope (and therefore need not be communicated between CDNs), while some of the Content Distribution Metadata may have an inter-CDN scope (and therefore needs to be communicated between CDNs).

CDNI Metadata: Content Distribution Metadata with inter-CDN scope. For example, CDNI Metadata may include geo-blocking information (i.e. information defining geographical areas where the content is to be made available or blocked), availability windows (i.e. information defining time windows during which the content is to be made available or blocked) and access control mechanisms to be enforced (e.g. URI signature validation). CDNI Metadata may also include information about desired distribution policy (e.g. prepositioned vs dynamic acquisition) and about where/how a CDN can acquire the content. CDNI Metadata may also include content management information (e.g. request for deletion of Content from Surrogates) across interconnected CDNs.

Dynamic content acquisition: Dynamic content acquisition is where a CDN acquires content from the content source in response to an End User requesting that content from the CDN. In the context of CDN Interconnection, dynamic acquisition means that a downstream CDN acquires the content from content sources (including upstream CDNs) at some point in time after a request for that content is delegated to the downstream CDN by an Upstream CDN (and that specific content is not yet available in the downstream CDN).

Dynamic CDNI metadata acquisition: In the context of CDN Interconnection, dynamic CDNI metadata acquisition means that a downstream CDN acquires CDNI metadata for content from the upstream CDN at some point in time after a request for that content is delegated to the downstream CDN by an Upstream CDN (and that specific CDNI metadata is not yet available in the downstream CDN). See also the definitions for downstream CDN and upstream CDN.

Pre-positioned content acquisition: Content Pre-positioning is where a CDN acquires content from the content source prior to, or independently of, any End User requesting that content from the CDN. In the context of CDN interconnection the Upstream CDN instructs the Downstream CDN to acquire the content from content sources (including upstream CDNs) in advance of or independent of any End User requesting it.

Pre-positioned CDNI Metadata acquisition: In the context of CDN Interconnection, CDNI Metadata pre-positioning is where the Downstream CDN acquires CDNI metadata for content prior to or independent of any End User requesting that content from the Downstream CDN.

End User (EU): The 'real' user of the system, typically a human but maybe some combination of hardware and/or software emulating a human (e.g. for automated quality monitoring etc.)

User Agent (UA): Software (or a combination of hardware and software) through which the End User interacts with a Content Service. The User Agent will communicate with a Content Service for the selection of content and one or more CDNs for the delivery of the Content. Such communication is not restricted to HTTP and may be via a variety of protocols. Examples of User Agents (non-exhaustive) are: Browsers, Set Top Boxes (STBs), dedicated content applications (e.g. media players), etc.

Network Service Provider (NSP): Provides network-based connectivity/ services to End Users.

Content Service Provider (CSP): Provides a Content Service to End Users (which they access via a User Agent). A CSP may own the Content made available as part of the Content Service, or may license content rights from another party.

Content Service: The service offered by a Content Service Provider. The Content Service encompasses the complete service which may be wider than just providing access to items of Content, e.g. the Content Service also includes any middleware, key distribution, program guide, etc. which may not require any direct interaction with the CDN, or CDNs, involved in the distribution and delivery of the content.

Content Distribution Network (CDN) / Content Delivery Network (CDN): Network infrastructure in which the network elements cooperate at layers 4 through layer 7 for more effective delivery of Content to User Agents. Typically a CDN consists of a Request Routing system, a Distribution System (that includes a set of Surrogates), a Logging

System and a CDN control system.

CDN Provider: The service provider who operates a CDN and offers a service of content delivery, typically used by a Content Service Provider or another CDN Provider. Note that a given entity may operate in more than one role. For example, a company may simultaneously operate as a Content Service Provider, a Network Service Provider and a CDN Provider.

CDN Interconnection (CDNI): A relationship between a pair of CDNs that enables one CDN to provide content delivery services on behalf of another CDN. A CDN Interconnection may be wholly or partially realized through a set of interfaces over which a pair of CDNs communicate with each other in order to achieve the delivery of content to User Agents by Surrogates in one CDN (the downstream CDN) on behalf of another CDN (the upstream CDN).

Authoritative CDN: A CDN which has a direct relationship with a CSP for the distribution & delivery of that CSP's content by the authoritative CDN or by downstream CDNs of the authoritative CDN.

Upstream CDN: For a given End User request, the CDN (within a pair of directly interconnected CDNs) that redirects the request to the other CDN.

Downstream CDN: For a given End User request, the CDN (within a pair of directly interconnected CDNs) to which the request is redirected by the other CDN (the Upstream CDN). Note that in the case of successive redirections (e.g. CDN1-->CDN2-->CDN3) a given CDN (e.g. CDN2) may act as the Downstream CDN for a redirection (e.g. CDN1-->CDN2) and as the Upstream CDN for the subsequent redirection of the same request (e.g. CDN2-->CDN3).

Over-the-top (OTT): A service, e.g. content delivery using a CDN, operated by a different operator than the NSP to which the users of that service are attached.

Surrogate: A device/function (often called a cache) that interacts with other elements of the CDN for the control and distribution of Content within the CDN and interacts with User Agents for the delivery of the Content. Typically, surrogates will cache requested content so that it can deliver the same content to a number of User Agents (and their End Users) avoiding the need for those requests to transit multiple times through the network core (i.e from the content origin to the surrogate).

Request Routing System: The function within a CDN responsible for receiving a content request from a User Agent, obtaining and

maintaining necessary information about a set of candidate surrogates or candidate CDNs, and for selecting and redirecting the user to the appropriate surrogate or CDN. To enable CDN Interconnection, the Request Routing System must also be capable of handling User Agent content requests passed to it by another CDN.

Distribution System: The function within a CDN responsible for distributing Content Distribution Metadata as well as the Content itself inside the CDN (e.g. down to the surrogates).

Delivery: The function within CDN surrogates responsible for delivering a piece of content to the User Agent. For example, delivery may be based on HTTP progressive download or HTTP adaptive streaming.

Logging System: The function within a CDN responsible for collecting the measurement and recording of distribution and delivery activities. The information recorded by the logging system may be used for various purposes including charging (e.g. of the CSP), analytics and monitoring.

Control System: The function within a CDN responsible for bootstrapping and controlling the other components of the CDN as well as for handling interactions with external systems (e.g. handling delivery service creation/update/removal requests, or specific service provisioning requests).

Quality of Experience (QoE): As defined in [Section 2.4 of \[RFC6390\]](#)

1.2. CDN Background

Readers are assumed to be familiar with the architecture, features and operation of CDNs. For readers less familiar with the operation of CDNs, the following resources may be useful:

- o [RFC 3040](#) [[RFC3040](#)] describes many of the component technologies that are used in the construction of a CDN.
- o Taxonomy [[TAXONOMY](#)] compares the architecture of a number of CDNs.
- o [RFC 3466](#) [[RFC3466](#)] and [RFC 3570](#) [[RFC3570](#)] are the output of the IETF Content Delivery Internetworking (CDI) working group which was closed in 2003.

Note: Some of the terms used in this document are similar to terms used the above referenced documents. When reading this document terms should be interpreted as having the definitions provided in [Section 1.1](#).

2. CDN Interconnection Use Cases

An increasing number of NSPs are deploying CDNs in order to deal cost-effectively with the growing usage of on-demand video services and other content delivery applications.

CDNs allow caching of content closer to the edge of a network so that a given item of content can be delivered by a CDN Surrogate (i.e. a cache) to multiple User Agents (and their End Users) without transiting multiple times through the network core (i.e from the content origin to the surrogate). This contributes to bandwidth cost reductions for the NSP and to improved quality of experience for the End Users. CDNs also enable replication of popular content across many surrogates, which enables content to be served to large numbers of User Agents concurrently. This also helps dealing with situations such as flash crowds and denial of service attacks.

The CDNs deployed by NSPs are not just restricted to the delivery of content to support the Network Service Provider's own 'walled garden' services, such as IP delivery of television services to Set Top Boxes, but are also used for delivery of content to other devices including PCs, tablets, mobile phones etc.

Some service providers operate over multiple geographies and federate multiple affiliate NSPs. These NSPs typically operate independent CDNs. As they evolve their services (e.g. for seamless support of content services to nomadic users across affiliate NSPs) there is a need for interconnection of these CDNs, that represents a first use case for CDNI. However there are no open specifications, nor common best practices, defining how to achieve such CDN interconnection.

CSPs have a desire to be able to get (some of) their content to very large numbers of End Users, who are often distributed across a number of geographies, while maintaining a high quality of experience, all without having to maintain direct business relationships with many different CDN Providers (or having to extend their own CDN to a large number of locations). Some NSPs are considering interconnecting their respective CDNs (as well as possibly over-the-top CDNs) so that this collective infrastructure can address the requirements of CSPs in a cost effective manner. This represents a second use case for CDNI. In particular, this would enable the CSPs to benefit from on-net delivery (i.e. within the Network Service Provider's own network/CDN footprint) whenever possible and off-net delivery otherwise, without requiring the CSPs to maintain direct business relationships with all the CDNs involved in the delivery. Again, CDN Providers (NSPs or over-the-top CDN operators) are faced with a lack of open specifications and best practices.

NSPs have often deployed CDNs as specialized cost-reduction projects within the context of a particular service or environment. Some NSPs operate separate CDNs for separate services. For example, there may be a CDN for managed IPTV service delivery, a CDN for web-TV delivery and a CDN for video delivery to Mobile terminals. As NSPs integrate their service portfolio, there is a need for interconnecting these CDNs, representing a third use case for CDNI. Again, NSPs face the problem of lack of open interfaces for CDN interconnection.

For operational reasons (e.g. disaster, flash crowd) or commercial reasons, an over-the-top CDN may elect to make use of another CDN (e.g. an NSP CDN with on-net Surrogates for a given footprint) for serving a subset of the user requests (e.g. requests from users attached to that NSP), which results in a fourth use case for CDNI because CDN Providers (over-the-top CDN Providers or NSPs) are faced with a lack of open specifications and best practices.

Use cases for CDN Interconnection are further discussed in [\[I-D.ietf-cdni-use-cases\]](#).

3. CDN Interconnection Model & Problem Area for IETF

This section discusses the problem area for the IETF work on CDN Interconnection.

Interconnecting CDNs involves interactions among multiple different functions and components that form each CDN. Only some of those require standardization.

Some NSPs have started to perform experiments to explore whether their CDN use cases can already be addressed with existing CDN implementations. One set of such experiments is documented in [\[I-D.bertrand-cdni-experiments\]](#). The conclusions of those experiments are that while some basic limited CDN Interconnection functionality can be achieved with existing CDN technology, the current lack of any standardized CDNI interfaces with the necessary level of functionality such as those discussed in this document is preventing the deployment of CDN Interconnection.

Listed below are the four interfaces required to interconnect a pair of CDNs and that constitute the problem space of CDN Interconnection along with the required functionality of each interface for which standards do not currently exist. As part of the development of the CDNI interfaces it will also be necessary to agree on common mechanisms for how to identify and name the data objects that are to be interchanged between interconnected CDNs.

The use of the term "interface" is meant to encompass the protocol over which CDNI data representations (e.g. CDNI Metadata objects) are exchanged as well as the specification of the data representations themselves (i.e. what properties/fields each object contains, its structure, etc.).


- o CDNI Control interface: This interface allows the "CDNI Control" system in interconnected CDNs to communicate. This interface may support the following:
 - * Allow bootstrapping of the other CDNI interfaces (e.g. interface address/URL discovery and establishment of security associations).
 - * Allow configuration of the other CDNI interfaces (e.g. Upstream CDN specifies information to be reported through the CDNI Logging interface).
 - * Allow the downstream CDN to communicate static (or fairly static) information about its delivery capabilities and policies.
 - * Allow bootstrapping of the interface between CDNs for content acquisition (even if that interface itself is outside the scope of the CDNI work).
 - * Allow an upstream CDN to initiate or request specific actions to be undertaken in the downstream CDN. For example, to allow an upstream CDN to initiate content or CDNI Metadata acquisition (pre-positioning) or to request the invalidation or purging of content files and/or CDNI Metadata in a downstream CDN.
- o CDNI Request Routing interface: This interface allows the Request Routing systems in interconnected CDNs to communicate to ensure that an End User request can be (re)directed from an upstream CDN to a surrogate in the downstream CDN, in particular where selection responsibilities may be split across CDNs (for example the upstream CDN may be responsible for selecting the downstream CDN while the downstream CDN may be responsible for selecting the actual surrogate within that downstream CDN). In particular, the functions of the CDN Request Routing interface may be divided as follows:
 - * A CDNI Request Routing Redirection interface which allows the upstream CDN to query the downstream CDN at request routing time before redirecting the request to the downstream CDN.
 - * A CDNI Footprint & Capabilities advertisement interface which allows the downstream CDN to provide to the upstream CDN (static or dynamic) information (e.g. resources, footprint, load) to facilitate selection of the downstream CDN by the upstream CDN request routing system when processing subsequent content requests from User Agents.

- o CDNI Metadata distribution interface: This interface allows the Distribution system in interconnected CDNs to communicate to ensure CDNI Metadata can be exchanged across CDNs. See [Section 1.1](#) for definition and examples of CDNI Metadata.
- o CDNI Logging interface: This interface allows the Logging system in interconnected CDNs to communicate the relevant activity logs in order to allow log consuming applications to operate in a multi-CDN environments. For example, an upstream CDN may collect delivery logs from a downstream CDN in order to perform consolidated charging of the CSP or for settlement purposes across CDNs. Similarly, an upstream CDN may collect delivery logs from a downstream CDN in order to provide consolidated reporting and monitoring to the CSP.

Note that the actual grouping of functionalities under these four interfaces is considered tentative at this stage and may be changed after further study (e.g. some subset of functionality be moved from one interface into another).

The above list covers a significant potential problem space, in part because in order to interconnect two CDNs there are several 'touch points' that require standardization. However, it is expected that the CDNI interfaces need not be defined from scratch and instead can very significantly reuse or leverage existing protocols; this is discussed further in [Section 4](#).

The interfaces that form the CDNI problem area are illustrated in Figure 1.



<==> interfaces inside the scope of CDNI
 ***** interfaces outside the scope of CDNI
 interfaces outside the scope of CDNI

Figure 1: A Model for the CDNI Problem Area

As illustrated in Figure 1, the acquisition of content between

interconnected CDNs is out of scope for CDNI, which deserves some additional explanation. The consequence of such a decision is that the CDNI problem space described in this document is focussed on only defining the control plane for CDNI; and the CDNI data plane (i.e. the acquisition & distribution of the actual content objects) is out of scope. The rationale for such a decision is that CDNs today typically already use standardized protocols such as HTTP, FTP, rsync, etc. to acquire content from their CSP customers and it is expected that the same protocols could be used for acquisition between interconnected CDNs. Therefore the problem of content acquisition is considered already solved and all that is required from specifications developed by the CDNI working group is to describe within the CDNI Metadata the parameters to use to retrieve the content for example the IP address/hostname to connect to, what protocol to use to retrieve the content, etc.

4. Scoping the CDNI Problem

This section outlines how the scope of work addressing the CDNI problem space can be constrained through reuse or leveraging of existing protocols to implement the CDNI interfaces. This discussion is not intended to pre-empt any working group decision as to the most appropriate protocols, technologies and solutions to select to realize the CDNI interfaces but is intended as an illustration of the fact that the CDNI interfaces need not be created in a vacuum and that reuse or leverage of existing protocols is likely possible.

The four CDNI interfaces (CDNI Control interface, CDNI Request Routing interface, CDNI Metadata interface, CDNI Logging interface) described in [Section 3](#) within the CDNI problem area are all control plane interfaces operating at the application layer (Layer 7 in the OSI network model). Firstly, since it is not expected that these interfaces would exhibit unique session, transport or network requirements as compared to the many other existing applications in the Internet, it is expected that the CDNI interfaces will be defined on top of existing session, transport and network protocols.

Secondly, although a new application protocol could be designed specifically for CDNI our analysis below shows that this is unnecessary and it is recommended that existing application protocols be reused or leveraged (HTTP [[RFC2616](#)], Atom Publishing Protocol [[RFC5023](#)], XMPP [[RFC6120](#)], for example) to realize the CDNI interfaces.

4.1. CDNI Request Routing Interface

The CDNI Request Routing interface enables a Request Routing function in an upstream CDN to query a Request Routing function in a downstream CDN to determine if the downstream CDN is able (and willing) to accept the delegated content request. It also allows the downstream CDN to control what should be returned to the User Agent in the redirection message by the upstream Request Routing function .

The CDNI Request Routing interface is therefore a fairly straightforward request/response interface and could be implemented over any number of request/response protocols. For example, it may be implemented as a WebService using one of the common WebServices methodologies (XML-RPC, HTTP query to a known URI, etc.). This removes the need for the CDNI working group to define a new protocol for the request/response element of the CDNI Request Routing interface.

Additionally, as discussed in [Section 3](#), the CDNI Request Routing interface is also expected to enable a downstream CDN to provide to the upstream CDN (static or dynamic) information (e.g. resources, footprint, load) to facilitate selection of the downstream CDN by the upstream CDN request routing system when processing subsequent content requests from User Agents. It is expected that such functionality of the CDNI request Routing could be specified by the CDNI working group with significant leveraging of existing IETF protocols supporting the dynamic distribution of reachability information (for example by leveraging existing routing protocols) or supporting application level queries for topological information (for example by leveraging ALTO [[RFC5693](#)]).

4.2. CDNI Metadata Interface

The CDNI Metadata interface enables the Distribution System in a downstream CDN to request CDNI Metadata from an upstream CDN so that the downstream CDN can properly process and respond to redirection requests received over the CDNI Request Routing interface and Content Requests received directly from User Agents.

The CDNI Metadata interface is therefore similar to the CDNI Request Routing interface because it is a request/response interface with the potential addition that CDNI Metadata search may have more complex semantics than a straightforward Request Routing redirection request. Therefore, like the CDNI Request Routing interface, the CDNI Metadata interface may be implemented as a WebService using one of the common WebServices methodologies (XML-RPC, HTTP query to a known URI, etc.) or possibly using other existing protocols such as XMPP [[RFC6120](#)]. This removes the need for the CDNI working group to define a new

protocol for the request/response element of the CDNI Metadata interface.

4.3. CDNI Logging Interface

The CDNI Logging interface enables details of content distribution and delivery activities to be exchanged between interconnected CDNs. For example the exchange of log records related to the delivery of content, similar to the log records recorded in a web server's access log.

Several protocols already exist that could potentially be used to exchange CDNI logs between interconnected CDNs including SNMP, syslog, ftp (and secure variants), HTTP POST, etc.

4.4. CDNI Control Interface

The CDNI Control interface allows the Control System in interconnected CDNs to communicate. The exact inter-CDN control functionality required to be supported by the CDNI Control interface is less well defined than the other three CDNI interfaces at this time.

It is expected that for the Control interface, as for the other CDNI Interfaces, existing protocols can be reused or leveraged.

5. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

6. Security Considerations

Distribution of content by a CDN comes with a range of security considerations such as how to enforce control of access to the content by end users in line with the CSP policy, or how to trust the logging information generated by the CDN for the purposes of charging the CSP. These security aspects are already dealt with by CDN Providers and CSPs today in the context of standalone CDNs. However, interconnection of CDNs introduces a new set of security considerations by extending the trust model to a chain of trust (i.e. the CSP "trusts" a CDN that "trusts" another CDN). The mechanisms used to mitigate these risks in multi-CDN environments may be similar to those used in the single CDN case, but their suitability in this

more complex environment must be validated.

The interconnection of CDNs may also introduce additional privacy considerations on top of those that apply to the single CDN case. In a multi-CDN environment, the different CDNs may reside in different legal regimes that require differing privacy requirements to be enforced. Such privacy requirements may impact the granularity of information that can be exchanged across the CDNI interfaces. For example the Logging System in a downstream CDN may need to apply some degree of anonymization, obfuscation or even the complete removal of some fields before exchanging log records containing details of End User deliveries with an upstream CDN.

Maintaining the security of the content itself, its associated metadata (including delivery policies) and the CDNs distributing and delivering it, are critical requirements for both CDN Providers and CSPs and the CDN Interconnection interfaces must provide sufficient mechanisms to maintain the security of the overall system of interconnected CDNs as well as the information (content, metadata, logs, etc) distributed and delivered through any set of interconnected CDNs.

6.1. Security of the CDNI Control interface

Information exchanged between interconnected CDNs over this interface is of a sensitive nature. A pair of CDNs use this interface to allow bootstrapping of all the other CDNI interfaces possibly including establishment of the mechanisms for securing these interfaces. Therefore, corruption of that interface may result in corruption of all other interfaces. Using this interface, an upstream CDN may pre-position or delete content or metadata in a downstream CDN and a downstream CDN may provide administrative information to an upstream CDN, etc. All of these operations require that the peer CDNs are appropriately authenticated and that the confidentiality and integrity of information flowing between them can be ensured.

6.2. Security of the CDNI Request Routing Interface

Appropriate levels of authentication and confidentiality must be used in this interface because it allows an upstream CDN to query the downstream CDN in order to redirect requests, and conversely, allows the downstream CDN to influence the upstream CDN's Request Routing function.

In the absence of appropriate security on this interface, a rogue upstream CDN could inundate downstream CDNs with bogus requests, or have the downstream CDN send the rogue upstream CDN private information. Also, a rogue downstream CDN could influence the

upstream CDN so the upstream CDN redirects requests to the rogue dCDN or another dCDN in order to, for example, attract additional delivery revenue.

6.3. Security of the CDNI Metadata interface

This interface allows a downstream CDN to request CDNI metadata from an upstream CDN, and therefore the upstream CDN must ensure that the former is appropriately authenticated before sending the data. Conversely, a downstream CDN must authenticate an upstream CDN before requesting metadata to insulate itself from poisoning by rogue upstream CDNs. The confidentiality and integrity of the information exchanged between the peers must be protected.

6.4. Security of the CDNI Logging interface

Logging data consists of potentially sensitive information (which end user accessed which media resource, IP addresses of end users, potential names and subscriber account information, etc.). Confidentiality of this information must be protected as log records are moved between CDNs. This information may also be sensitive from the viewpoint that it can be the basis for charging across CDNs. Therefore, appropriate levels of protection are needed against corruption, duplication and loss of this information.

7. Acknowledgements

The authors would like to thank Andre Beck, Gilles Bertrand, Mark Carlson, Bruce Davie, David Ferguson, Yiu Lee, Kent Leung, Will Li, Kevin Ma, Julien Maisonneuve, Guy Meador, Larry Peterson, Emile Stephan, Oskar van Deventer, Mahesh Viveganandhan and Richard Woundy for their review comments and contributions to the text.

8. References

8.1. Normative References

8.2. Informative References

[3GP-DASH]

"Transparent end-to-end Packet-switched Streaming Service (PSS); Progressive Download and Dynamic Adaptive Streaming over HTTP (3GP-DASH)

<http://www.3gpp.org/ftp/Specs/html-info/26247.htm>".

[ALTO-Charter]

"IETF ALTO WG Charter
(<http://datatracker.ietf.org/wg/alto/charter/>)".

[ATIS] "ATIS (<http://www.atis.org/>)".

[ATIS-COD]

"ATIS IIF: IPTV Content on Demand Service, January 2011 (http://www.atis.org/iif/_Com/Docs/Task_Forces/ARCH/ATIS-0800042.pdf)".

[CDI-Charter]

"IETF CDI WG Charter
(<http://www.ietf.org/wg/concluded/cdi>)".

[CableLabs]

"CableLabs (<http://www.cablelabs.com/about/>)".

[CableLabs-Metadata]

"CableLabs VoD Metadata Project Primer
(<http://www.cablelabs.com/projects/metadata/primer/>)".

[DECADE-Charter]

"IETF DECADE WG Charter
(<http://datatracker.ietf.org/wg/decade/charter/>)".

[I-D.bertrand-cdni-experiments]

Faucheur, F. and L. Peterson, "Content Distribution Network Interconnection (CDNI) Experiments", [draft-bertrand-cdni-experiments-02](#) (work in progress), February 2012.

[I-D.ietf-cdni-use-cases]

Bertrand, G., Emile, S., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", [draft-ietf-cdni-use-cases-08](#) (work in progress), June 2012.

[I-D.jenkins-alto-cdn-use-cases]

Niven-Jenkins, B., Watson, G., Bitar, N., Medved, J., and S. Previdi, "Use Cases for ALTO within CDNs", [draft-jenkins-alto-cdn-use-cases-03](#) (work in progress), June 2012.

[MPEG-DASH]

"Information technology - MPEG systems technologies - Part 6: Dynamic adaptive streaming over HTTP (DASH), (DIS version), February 2011
<http://mpeg.chiariglione.org/>

working_documents.htm#MPEG-B".

[OIPF-Overview]

"OIPF Release 2 Specification Volume 1 - Overview",
September 2010.

[P2PRG-CDNI]

Davie, B. and F. Le Faucheur, "Interconnecting CDNs aka
"Peering Peer-to-Peer"
(<http://www.ietf.org/proceedings/77/slides/P2PRG-2.pdf>)",
March 2010.

[PPSP-Charter]

"IETF PPSP WG Charter
(<http://datatracker.ietf.org/wg/ppsp/charter/>)".

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H.,
Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext
Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.

[RFC3040] Cooper, I., Melve, I., and G. Tomlinson, "Internet Web
Replication and Caching Taxonomy", [RFC 3040](#), January 2001.

[RFC3466] Day, M., Cain, B., Tomlinson, G., and P. Rzewski, "A Model
for Content Internetworking (CDI)", [RFC 3466](#),
February 2003.

[RFC3568] Barbir, A., Cain, B., Nair, R., and O. Spatscheck, "Known
Content Network (CN) Request-Routing Mechanisms",
[RFC 3568](#), July 2003.

[RFC3570] Rzewski, P., Day, M., and D. Gilletti, "Content
Internetworking (CDI) Scenarios", [RFC 3570](#), July 2003.

[RFC5023] Gregorio, J. and B. de h0ra, "The Atom Publishing
Protocol", [RFC 5023](#), October 2007.

[RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic
Optimization (ALTO) Problem Statement", [RFC 5693](#),
October 2009.

[RFC6120] Saint-Andre, P., "Extensible Messaging and Presence
Protocol (XMPP): Core", [RFC 6120](#), March 2011.

[RFC6390] Clark, A. and B. Claise, "Guidelines for Considering New
Performance Metric Development", [BCP 170](#), [RFC 6390](#),
October 2011.

[SNIA-CDMI]

"SNIA CDMI (http://www.snia.org/tech_activities/standards/curr_standards/cdmi)".

[TAXONOMY]

Pathan, A., "A Taxonomy and Survey of Content Delivery Networks (<http://www.gridbus.org/reports/CDN-Taxonomy.pdf>)", 2007.

[Y.1910] "ITU-T Recommendation Y.1910 "IPTV functional architecture"", September 2008.

[Y.2019] "ITU-T Recommendation Y.2019 "Content delivery functional architecture in NGN"", September 2010.

Appendix A. Design considerations for realizing the CDNI Interfaces

This section expands on how CDNI interfaces can reuse and leverage existing protocols before describing each CDNI interface individually and highlighting example candidate protocols that could be considered for reuse or leveraging to implement the CDNI interfaces.

A.1. CDNI Request Routing Interface

The CDNI Request Routing interface enables a Request Routing function in an upstream CDN to query a Request Routing function in a downstream CDN to determine if the downstream CDN is able (and willing) to accept the delegated content request and to allow the downstream CDN to control what the upstream Request Routing function should return to the User Agent in the redirection message.

Therefore, the CDNI Request Routing interface needs to offer a mechanism for an upstream CDN to issue a "Redirection Request" to a downstream CDN. The Request Routing interface needs to be able to support scenarios where the initial User Agent request to the upstream CDN is received over DNS as well as over a content specific application protocol (e.g. HTTP, RTSP, RTMP, etc.).

Therefore a Redirection Request is expected to contain information such as:

- o The protocol (e.g. DNS, HTTP) over which the upstream CDN received the initial User Agent request.
- o Additional details of the User Agent request that are required to perform effective Request Routing by the Downstream CDN. For DNS this would typically be the IP address of the DNS resolver making the request on behalf of the User Agent. For requests received

over content specific application protocols the Redirection Request could contain significantly more information related to the original User Agent request but at a minimum is expected to include the User Agent's IP address, the equivalent of the HTTP Host header and the equivalent of the HTTP abs_path defined in [\[RFC2616\]](#).

It should be noted that, the CDNI architecture needs to consider that a downstream CDN may receive requests from User Agents without first receiving a Redirection Request from an upstream CDN for the corresponding User Agent request, for example because:

- o User Agents (or DNS resolvers) may cache DNS or application responses from Request Routers.
- o Responses to Redirection Requests over the Request Routing interface may be cacheable.
- o Some CDNs may rely on simple coarse policies, e.g. CDN B agrees to always serve CDN A's delegated redirection requests, in which case the necessary redirection details are exchanged out of band (of the CDNI interfaces), e.g. configured.

On receiving a Redirection Request, the downstream CDN will use the information provided in the request to determine if it is able (and willing) to accept the delegated content request and needs to return the result of its decision to the upstream CDN.

Thus, a Redirection Response from the downstream CDN is expected to contain information such as:

- o Status code indicating acceptance or rejection (possibly with accompanying reasons).
- o Information to allow redirection by the Upstream CDN. In the case of DNS-based request routing, this is expected to include the equivalent of a DNS record(s) (e.g. a CNAME) that the upstream CDN should return to the requesting DNS resolver. In the case of application based request routing, this is expected to include the information necessary to construct the application specific redirection response(s) to return to the requesting User Agent. For HTTP requests from User Agents this could include a URI that the upstream CDN could return in a HTTP 3xx response.

The CDNI Request Routing interface is therefore a fairly straightforward request/response interface and could be implemented over any number of request/response protocols. For example, it may be implemented as a WebService using one of the common WebServices methodologies (XML-RPC, HTTP query to a known URI, etc.). This removes the need for the CDNI working group to define a new protocol for the request/response element of the CDNI Request Routing

interface. Thus, the CDNI working group would be left only with the task of specifying:

- o The recommended request/response protocol to use along with any additional semantics and procedures that are specific to the CDNI Request Routing interface (e.g. handling of malformed requests/responses).
- o The syntax (i.e representation/encoding) of the redirection requests and responses.
- o The semantics (i.e. meaning and expected contents) of the redirection requests and responses.

Additionally, as discussed in [Section 3](#), the CDNI Request Routing interface is also expected to enable a downstream CDN to provide to the upstream CDN (static or dynamic) information (e.g. resources, footprint, load) to facilitate selection of the downstream CDN by the upstream CDN request routing system when processing subsequent content requests from User Agents. It is expected that such functionality of the CDNI request Routing could be specified by the CDNI working group with significant leveraging of existing IETF protocols supporting the dynamic distribution of reachability information (for example by leveraging existing routing protocols) or supporting application level queries for topological information (for example by leveraging ALTO).

[A.2.](#) CDNI Metadata Interface

The CDNI Metadata interface enables the Distribution System in a downstream CDN to obtain CDNI Metadata from an upstream CDN so that the downstream CDN can properly process and respond to:

- o Redirection Requests received over the CDNI Request Routing interface.
- o Content Requests received directly from User Agents.

The CDNI Metadata interface needs to offer a mechanism for an Upstream CDN to:

- o Distribute/update/remove CDNI Metadata to a Downstream CDN.

and/or to allow a downstream CDN to:

- o Make direct requests for CDNI Metadata objects
- o Make recursive requests for CDNI metadata, for example to enable a downstream CDN to walk down a tree of objects with inter-object relationships.

The CDNI Metadata interface is therefore similar to the CDNI Request

Routing interface because it is a request/response interface with the potential addition that CDNI Metadata search may have more complex semantics than a straightforward Request Routing redirection request. Therefore, like the CDNI Request Routing interface, the CDNI Metadata interface may be implemented as a WebService using one of the common WebServices methodologies (XML-RPC, HTTP query to a known URI, etc.) or possibly using other existing protocols such as XMPP [[RFC6120](#)]. This removes the need for the CDNI working group to define a new protocol for the request/response element of the CDNI Metadata interface.

Thus, the CDNI working group would be left only with the task of specifying:

- o The recommended request/response protocol to use along with any additional semantics that are specific to the CDNI Metadata interface (e.g. handling of malformed requests/responses).
- o The syntax (i.e representation/encoding) of the CDNI Metadata objects that will be exchanged over the interface.
- o The semantics (i.e. meaning and expected contents) of the individual properties of a Metadata object.
- o How the relationships between different CDNI Metadata objects are represented.

[A.3.](#) CDNI Logging Interface

The CDNI Logging interface enables details of logs or events to be exchanged between interconnected CDNs, where events could be:

- o Log records related to the delivery of content (similar to the log records recorded in a web server's access log).
- o Real-time or near-real time events before, during or after content delivery, e.g. content delivery interruption
- o Operations and diagnostic messages.

Within CDNs today, logs and events are used for a variety of purposes in addition to real-time and non real-time diagnostics and auditing by the CDN Provider and its customers. Specifically CDNs use logs to generate Call Data Records (CDRs) for passing to billing and payment systems and to real-time (and near real-time) analytics systems. Such applications place requirements on the CDNI Logging interface to support guaranteed and timely delivery of log messages between interconnected CDNs. It may also be necessary to be able to prove the integrity of received log messages.

Several protocols already exist that could potentially be used to exchange CDNI logs between interconnected CDNs including SNMP Traps, syslog, ftp, HTTP POST, etc. although it is likely that some of the

candidate protocols may not be well suited to meet all the requirements of CDNI. For example SNMP traps pose scalability concerns and SNMP does not support guaranteed delivery of Traps and therefore could result in log records being lost and the consequent CDRs and billing records for that content delivery not being produced as well as that content delivery being invisible to any analytics platforms.

Although it is not necessary to define a new protocol for exchanging logs across the CDNI Logging interface, the CDNI working group would still need to specify:

- o The recommended protocol to use.
- o A default set of log fields and their syntax & semantics. Today there is no standard set of common log fields across different content delivery protocols and in some cases there is not even a standard set of log field names and values for different implementations of the same delivery protocol.
- o A default set of events that trigger logs to be generated.

A.4. CDNI Control Interface

The CDNI Control interface allows the Control System in interconnected CDNs to communicate. The exact inter-CDN control functionality required to be supported by the CDNI Control interface is less well defined than the other three CDNI interfaces at this time.

However, as discussed in [Section 3](#), the CDNI Control interface may be required to support functionality similar to the following:

- o Allow an upstream CDN and downstream CDN to establish, update or terminate their CDNI interconnection.
- o Allow bootstrapping of the other CDNI interfaces (e.g. protocol address discovery and establishment of security associations).
- o Allow configuration of the other CDNI interfaces (e.g. Upstream CDN specifies information to be reported through the CDNI Logging interface).
- o Allow the downstream CDN to communicate static information about its delivery capabilities, resources and policies.
- o Allow bootstrapping of the interface between CDNs for content acquisition (even if that interface itself is outside the scope of the CDNI work).

It is expected that for the Control interface also, existing protocols can be reused or leveraged. Those will be considered once the requirements for the Control interface have been refined.

Appendix B. Additional Material

This section records related information that was produced as part of defining the CDNI problem statement.

B.1. Non-Goals for IETF

Listed below are aspects of content delivery that the authors propose be kept outside of the scope of the CDNI working group:

- o The interface between Content Service Provider and the Authoritative CDN (i.e. the upstream CDN contracted by the CSP for delivery by this CDN or by its downstream CDNs).
- o The delivery interface between the delivering CDN surrogate and the User Agent, such as streaming protocols.
- o The request interface between the User Agent and the request-routing system of a given CDN. Existing IETF protocols (e.g. HTTP, RTSP, DNS) are commonly used by User Agents to request content from a CDN and by CDN request routing systems to redirect the User Agent requests. The CDNI working group need not define new protocols for this purpose. Note however, that the CDNI control plane interface may indirectly affect some of the information exchanged through the request interface (e.g. URI).
- o The content acquisition interface between CDNs (i.e. the data plane interface for actual delivery of a piece of content from one CDN to the other). This is expected to use existing protocols such as HTTP or protocols defined in other forums for content acquisition between an origin server and a CDN (e.g. HTTP-based C2 reference point of ATIS IIF CoD). The CDN Interconnection problem space described in this document may therefore only concern itself with the agreement/negotiation aspects of which content acquisition protocol is to be used between two interconnected CDNs in view of facilitating interoperability.
- o End User/User Agent Authentication. End User/User Agent authentication and authorization are the responsibility of the Content Service Provider.
- o Content preparation, including encoding and transcoding. The CDNI architecture aims at allowing distribution across interconnected CDNs of content treated as opaque objects. Interpretation and processing of the objects, as well as optimized delivery of these objects by the surrogate to the End User are outside the scope of CDNI.
- o Digital Rights Management (DRM). DRM is an end-to-end issue between a content protection system and the User Agent.
- o Applications consuming CDNI logs (e.g. charging, analytics, reporting,...).
- o Internal CDN interfaces & protocols (i.e. interfaces & protocols within one CDN).

- o Scalability of individual CDNs. While scalability of the CDNI interfaces/approach is in scope, how an individual CDN scales is out of scope.
- o Actual algorithms for selection of CDNs or Surrogates by Request Routing systems (however, some specific parameters required as input to these algorithms may be in scope when they need to be communicated across CDNs).
- o Surrogate algorithms. For example caching algorithms and content acquisition methods are outside the scope of the CDNI work. Content management (e.g. Content Deletion) as it relates to CDNI content management policies, is in scope but the internal algorithms used by a cache to determine when to no longer cache an item of Content (in the absence of any specific metadata to the contrary) is out of scope.
- o Element management interfaces.
- o Commercial, business and legal aspects related to the interconnections of CDNs.

B.2. Relationship to relevant IETF Working Groups & IRTF Reserach Groups

B.2.1. ALTO WG

As stated in the ALTO Working Group charter [[ALTO-Charter](#)]:

"The Working Group will design and specify an Application-Layer Traffic Optimization (ALTO) service that will provide applications with information to perform better-than-random initial peer selection. ALTO services may take different approaches at balancing factors such as maximum bandwidth, minimum cross-domain traffic, lowest cost to the user, etc. The working group will consider the needs of BitTorrent, tracker-less P2P, and other applications, such as content delivery networks (CDN) and mirror selection."

In particular, the ALTO service can be used by a CDN Request Routing system to improve its selection of a CDN surrogate to serve a particular User Agent request (or to serve a request from another surrogate). [[I-D.jenkins-alto-cdn-use-cases](#)] describes a number of use cases for a CDN to be able to obtain network topology and cost information from an ALTO server(s) and discusses how CDN Request Routing could be used as an integration point of ALTO into CDNs. It is possible that the ALTO service could be used in the same manner in a multi-CDN environment based on CDN Interconnection. For example, an upstream CDN may take advantage of the ALTO service in its decision for selecting a downstream CDN to which a user request should be delegated.

However, the current work of ALTO is complementary to and does not

overlap with the work described in this document because the integration between ALTO and a CDN is an internal decision for a specific CDN and is therefore out of scope for the CDNI working group. One area for further study is whether additional information should be provided by an ALTO service to facilitate CDNI CDN selection.

B.2.2. DECADE WG

The DECADE Working Group [[DECADE-Charter](#)] is addressing the problem of reducing traffic on the last-mile uplink, as well as backbone and transit links caused by P2P streaming and file sharing applications. It addresses the problem by enabling an application endpoint to make content available from an in-network storage service and by enabling other application endpoints to retrieve the content from there.

Exchanging data through the in-network storage service in this manner, instead of through direct communication, provides significant gain where:

- o The network capacity/bandwidth from in-network storage service to application endpoint significantly exceeds the capacity/bandwidth from application endpoint to application endpoint (e.g. because of an end-user uplink bottleneck); and
- o Where the content is to be accessed by multiple instances of application endpoints (e.g. as is typically the case for P2P applications).

While, as is the case for any other data distribution application, the DECADE architecture and mechanisms could potentially be used for exchange of CDNI control plane information via an in-network-storage service (as opposed to directly between the entities terminating the CDNI interfaces in the neighbor CDNs), we observe that:

- o CDNI would operate as a "Content Distribution Application" from the DECADE viewpoint (i.e. would operate on top of DECADE).
- o There does not seem to be obvious benefits in integrating the DECADE control plane responsible for signaling information relating to control of the in-network storage service itself, and the CDNI control plane responsible for application-specific CDNI interactions (such as exchange of CDNI metadata, CDNI request redirection, transfer of CDNI logging information).
- o There would typically be limited benefits in making use of a DECADE in-network storage service because the CDNI interfaces are expected to be terminated by a very small number of CDNI clients (if not one) in each CDN, and the CDNI clients are expected to benefit from high bandwidth/capacity when communicating directly to each other (at least as high as if they were communicating via

an in-network storage server).

The DECADE in-network storage architecture and mechanisms may theoretically be used for the acquisition of the content objects themselves between interconnected CDNs. It is not expected that this would have obvious benefits in typical situations where a content object is acquired only once from an Upstream CDN to a Downstream CDN (and then distributed as needed inside the Downstream CDN). But it might have benefits in some particular situations. Since the acquisition protocol between CDNs is outside the scope of the CDNI work, this question is left for further study.

The DECADE in-network storage architecture and mechanisms may potentially also be used within a given CDN for the distribution of the content objects themselves among surrogates of that CDN. Since the CDNI work does not concern itself with operation within a CDN, this question is left for further study.

Therefore, the work of DECADE may be complementary to but does not overlap with the CDNI work described in this document.

B.2.3. PPSP WG

As stated in the PPSP Working Group charter [[PPSP-Charter](#)]:

"The Peer-to-Peer Streaming Protocol (PPSP) working group develops two signaling and control protocols for a peer-to-peer (P2P) streaming system for transmitting live and time-shifted media content with near real-time delivery requirements." and "The PPSP working group designs a protocol for signaling and control between trackers and peers (the PPSP "tracker protocol") and a signaling and control protocol for communication among the peers (the PPSP "peer protocol"). The two protocols enable peers to receive streaming data within the time constraints required by specific content items."

Therefore PPSP is concerned with the distribution of the streamed content itself along with the necessary signaling and control required to distribute the content. As such, it could potentially be used for the acquisition of streamed content across interconnected CDNs. But since the acquisition protocol is outside the scope of the work proposed for CDNI, we leave this for further study. Also, because of its streaming nature, PPSP is not seen as applicable to the distribution and control of the CDNI control plane and CDNI data representations.

Therefore, the work of PPSP may be complementary to but does not overlap with the work described in this document for CDNI.

B.2.4. IRTF P2P Research Group

Some information on CDN interconnection motivations and technical issues were presented in the P2P RG at IETF 77. The presentation can be found in [[P2PRG-CDNI](#)].

Appendix C. Additional Material

Note to RFC Editor: This appendix is to be removed on publication as an RFC.

C.1. Related standardization activities

There are a number of other standards bodies and industry forums that are working in areas related to CDNs, and in some cases related to CDNI. This section outlines any potential overlap with the work of the CDNI working group and any component that could potentially be reused to realize the CDNI interfaces.

A number of standards bodies have produced specifications related to CDNs, for example:

- o ETSI TISPAN (Telecommunications and Internet converged Services and Protocols for Advanced Networking) has a series of specifications focusing on CDNs.
- o The Open IPTV Forum (OIPF) and ATIS IPTV Interoperability Forum (IIF) specify the architecture and the protocols of an IPTV solution. Although OIPF and ATIS specifications include the interaction with a CDN, the CDN specifications are coupled with their IPTV specifications and do not cover interconnection of CDNs.
- o ATIS Cloud Services Forum (CSF) has started investigating interconnection of CDNs. The ATIS CSF focuses on defining use cases and requirements for such CDN interconnection, which are expected to be considered as input into the work of the CDNI working group. At the time of writing this document, ATIS CSF is not specifying the corresponding protocols or interfaces and is expected to leverage the work of the IETF CDNI working group for those.
- o CableLabs, SNIA and ITU have developed (or are working on) definitions for content related metadata and specifications for its distribution. However, they do not include metadata specific to the distribution of content within a CDN or between interconnected CDNs.
- o IETF CDI working group (now concluded) touched on the same problem space as the present document. However, in accordance with its initial charter, the CDI working group did not define any

protocols or interfaces to actually enable CDN Interconnection and at that time (2003) there was not enough industry interest and real life requirements to justify rechartering the working group to conduct the corresponding protocol work.

Although some of the specifications describe multi-CDN cooperation or include reference points for interconnecting CDNs, none of them specify in sufficient detail all the CDNI interfaces and CDNI Metadata representations required to enable even a base level of CDN Interconnection functionality to be implemented.

C.1.1. IETF CDI Working Group (Concluded)

The Content Distribution Internetworking (CDI) Working Group was formed in the IETF following a BoF in December 2000 and closed in mid 2003.

For convenience, here is an extract from the CDI working group charter [[CDI-Charter](#)]:

"

- o The goal of this working group is to define protocols to allow the interoperation of separately-administered content networks.
- o A content network is an architecture of network elements, arranged for efficient delivery of digital content. Such content includes, but is not limited to, web pages and images delivered via HTTP, and streaming or continuous media which are controlled by RTSP.
- o The working group will first define requirements for three modes of content internetworking: interoperation of request-routing systems, interoperation of distribution systems, and interoperation of accounting systems. These requirements are intended to lead to a follow-on effort to define protocols for interoperation of these systems.
- o In its initial form, the working group is not chartered to deliver those protocols [...]

"

Thus, the CDI working group touched on the same problem space as the present document.

The CDI working group published 3 Informational RFCs:

- o [RFC 3466](#) [[RFC3466](#)] - "A Model for Content Internetworking (CDI)".
- o [RFC 3568](#) [[RFC3568](#)] - "Known Content Network (CN) Request-Routing Mechanisms".

- o [RFC 3570](#) [[RFC3570](#)] - "Content Internetworking (CDI) Scenarios".

[C.1.2.](#) 3GPP

3GPP was the first organization that released a specification related to adaptive streaming over HTTP. 3GPP Release 9 specification on adaptive HTTP streaming was published in March 2010, and there have been some bug fixes on this specification since the publication. In addition, 3GPP has produced an extended version for Release 10, which was published in 2011. This release will include a number of clarifications, improvements and new features.

[3GP-DASH] is defined as a general framework independent of the data encapsulation format. It has support for fast initial startup and seeking, adaptive bitrate switching, re-use of HTTP origin and cache servers, re-use of existing media playout engines, on-demand, live and time-shifted delivery. It specifies syntax and semantics of Media Presentation Description (MPD), format of segments and delivery protocol for segments. It does not specify content provisioning, client behavior or transport of MPD.

The content retrieved by a client using [[3GP-DASH](#)] adaptive streaming could be obtained from a CDN but this is not discussed or specified in the 3GPP specifications as it is transparent to [[3GP-DASH](#)] operations. Similarly, it is expected that [[3GP-DASH](#)] can be used transparently from the CDNs as a delivery protocol (between the delivering CDN surrogate and the User Agent) in a CDN Interconnection environment. [[3GP-DASH](#)] could also be a candidate for content acquisition between CDNs in a CDN Interconnection environment.

[C.1.3.](#) ISO MPEG

Within ISO MPEG, the Dynamic Adaptive Streaming over HTTP (DASH) ad-hoc group adopted the 3GPP Release 9 [[3GP-DASH](#)] specification as a starting point and has made some improvements and extensions. Similar to 3GPP SA4, the MPEG DASH ad-hoc group has been working on standardizing the manifest file and the delivery format. Additionally, the MPEG DASH ad-hoc group has also been working on the use of MPEG-2 Transport Streams as a media format, conversion from/to existing file formats, common encryption, and so on. The MPEG DASH specification could also be a candidate for delivery to the User Agent and for content acquisition between CDNs in a CDN Interconnection environment. The Draft International Standard (DIS) version [[MPEG-DASH](#)] is currently publicly available since early February 2011.

In the 95th MPEG meeting in January 2011, the DASH ad-hoc group decided to start a new evaluation experiment called "CDN-EE". The

goals are to understand the requirements for MPEG DASH to better support CDN-based delivery, and to provide a guidelines document for CDN operators to better support MPEG DASH streaming services. The ongoing work is still very preliminary and does not currently target looking into CDN Interconnection use cases.

C.1.4. ATIS IIF

ATIS ([\[ATIS\]](#)) IIF is the IPTV Interoperability Forum (within ATIS) that develops requirements, standards, and specifications for IPTV.

ATIS IIF is developing the "IPTV Content on Demand (CoD) Service" specification. This includes use of a CDN (referred to in ATIS IIF CoD as the "Content Distribution and Delivery Functions") for support of a Content on Demand (CoD) Service as part of a broader IPTV service. However, this only covers the case of a managed IPTV service (in particular where the CDN is administered by the service provider) and does not cover the use, or interconnection, of multiple CDNs.

C.1.5. CableLabs

"Founded in 1988 by cable operating companies, Cable Television Laboratories, Inc. (CableLabs) is a non-profit research and development consortium that is dedicated to pursuing new cable telecommunications technologies and to helping its cable operator members integrate those technical advancements into their business objectives." [\[CableLabs\]](#)

CableLabs has defined specifications for CoD Content Metadata as part of its VOD Metadata project.

C.1.6. ETSI MCD

ETSI MCD (Media Content Distribution) is the ETSI technical committee "in charge of guiding and coordinating standardization work aiming at the successful overall development of multimedia systems (television and communication) responding to the present and future market requests on media content distribution".

MCD created a specific work item on interconnection of heterogeneous CDNs ("CDN Interconnection, use cases and requirements") in March 2010. MCD very recently created a working group to progress this work item. However, no protocol level work has yet started in MCD for CDN Interconnection.

C.1.7. ETSI TISPAN

ETSI TISPAN has published two sets of IPTV specifications, one of which is based on IMS. In addition, TISPAN has published a CDN architecture supporting delivery of various content services such as time-shifted TV and VoD to TISPAN devices (UEs) or regular PCs. The use cases allow for hierarchically and geographically distributed CDN scenarios, along with multi-CDN cooperation. As a result, the architecture contains reference points to support interconnection of other TISPAN CDNs. The protocol definition phase for the corresponding CDN architecture was kicked-off at the end of 2010 as is still in progress. In line with its long history of leveraging IETF protocols, ETSI could potentially leverage CDNI interfaces developed in the IETF for their related protocol level work on interconnections of CDNs.

C.1.8. ITU-T

SG13 is developing standards related to the support of IPTV services (i.e.. multimedia services such as television/VoD/audio/text/graphics/data delivered over IP-based managed networks).

ITU-T Recommendation Y.1910 [[Y.1910](#)] provides the description of the IPTV functional architecture. This architecture includes functions and interfaces for the distribution and delivery of content. This architecture is aligned with the ATIS IIF architecture.

Based upon ITU-T Rec. Y.1910, ITU-T Rec. Y.2019 [[Y.2019](#)] describes in more detail the content delivery functional architecture. This architecture allows CDN Interconnection: some interfaces (such as D3, D4) at the control level allow relationships between different CDNs, in the same domain or in different domains. Generic procedures are described, but the choice of the protocols is open.

C.1.9. Open IPTV Forum (OIPF)

The Open IPTV Forum has developed an end-to-end solution to allow any OIPF terminal to access enriched and personalized IPTV services either in a managed or a non-managed network[OIPF-Overview]. Some OIPF services (such as Network PVR) may be hosted in a CDN.

To that end, the Open IPTV Forum specification is made of 5 parts:

- o Media Formats including HTTP Adaptive Streaming
- o Content Metadata
- o Protocols

- o Terminal (Declarative or Procedural Application Environment)
- o Authentication, Content Protection and Service Protection

C.1.10. TV-Anytime Forum

Version 1 of the TV-Anytime Forum specifications were published as ETSI TS 102 822-1 through ETSI TS 102 822-7 "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime)". It includes the specification of content metadata in XML schemas (ETSI TS 102 822-3) which define technical parameters for the description of CoD and Live contents. The specification is referenced by DVB and OIPF.

The TV-anytime Forum was closed in 2005.

C.1.11. SNIA

The Storage Networking Industry Association (SNIA) is an association of producers and consumers of storage networking products whose goal is to further storage networking technology and applications.

SNIA has published the Cloud Data Management Interface (CDMI) standard ([[SNIA-CDMI](#)]).

"The Cloud Data Management Interface defines the functional interface that applications will use to create, retrieve, update and delete data elements from the Cloud. As part of this interface the client will be able to discover the capabilities of the cloud storage offering and use this interface to manage containers and the data that is placed in them. In addition, metadata can be set on containers and their contained data elements through this interface."

C.1.12. Summary of existing standardization work

The following sections will summarize the existing work of the standard bodies listed earlier against the CDNI problem space. [Appendix C.1.12.1](#) summarizes existing interfaces that could be leveraged for content acquisition between CDNs and [Appendix C.1.12.2](#) summarizes existing metadata specifications that may be applicable to CDNI. To date we are not aware of any standardization activities in the areas of the remaining CDNI interfaces (CDNI Request Routing, CDNI Control and CDNI Logging).

C.1.12.1. Content Acquisition across CDNs and Delivery to End User (Data plane)

A number of standards bodies have completed work in the areas of content acquisition interface between a CSP and a CDN, as well as as

on the delivery interface between the surrogate and the User Agent. Some of this work is summarized below.

TISPAN, OIPF and ATIS have specified IPTV and/or Content on Demand (CoD) services, including the data plane aspects (typically different flavors of RTP/RTCP and HTTP) to obtain content and deliver it to User Agents. For example, :

- o The OIPF data plane includes both RTP and HTTP flavors (HTTP progressive download, HTTP Adaptive streaming [[3GP-DASH](#)]).
- o The ATIS IIF specification "IPTV Content on Demand (CoD) Service" [[ATIS-COD](#)] defines a reference point (C2) and the corresponding HTTP-based data plane protocol for content acquisition between an authoritative origin server and the CDN.

While these protocols have not been explicitly specified for content acquisition across CDNs, they are suitable (in addition to others such as standard HTTP) for content acquisition between CDNs in a CDN Interconnection environment. Therefore for the purpose of the CDNI working group there are already multiple existing data plane protocols that can be used for content acquisition across CDNs.

Similarly, there are multiple existing standards (e.g. the OIPF data plane mentioned above, HTTP adaptive streaming [[3GP-DASH](#)]) or public specifications (e.g. vendor specific HTTP Adaptive streaming specifications) so that content delivery can be considered already solved (or at least sufficiently addressed in other forums).

Thus, specification of the content acquisition interface between CDNs and the delivery interface between the surrogate and the User Agent are out of scope for the CDNI working group. The CDNI working group may only concern itself with the negotiation/selection aspects of the acquisition protocol to be used in a CDN interconnect scenario.

C.1.12.2. CDNI Metadata

CableLabs, ITU, OIPF and TV-Anytime have work items dedicated to the specification of content metadata:

- o CableLabs has defined specifications for CoD Content Metadata as part of its VOD Metadata project. "The VOD Metadata project is a cable television industry and cross-industry-wide effort to specify the metadata and interfaces for distribution of video-on-demand (VOD) material from multiple content providers to cable operators." [[CableLabs-Metadate](#)]. However, while the CableLabs work specifies an interface between a content provider and a service provider running a CDN, it does not include an interface that could be used between CDNs.

- o ITU Study Group 16 has started work on a number of draft Recommendations (H.IPTV-CPMD, H.IPTV-CPMD, HSTP.IPTV-CMA, HSTP.IPTV-UMCI) specifying metadata for content distribution in IPTV services.
- o An Open IPTV Terminal receives the technical description of the content distribution from the OIPF IPTV platform before receiving any content. The Content distribution metadata is sent in the format of a TV-Anytime XSD including tags to describes the location and program type (on demand or Live) as well as describing the time availability of the on demand and live content.

However the specifications outlined above do not include metadata specific to the distribution of content within a CDN or between interconnected CDNs, for example geo-blocking information, availability windows, access control mechanisms to be enforced by the surrogate, how to map an incoming content request to a file on the origin server or acquire it from the upstream CDN etc.

The CDMI standard ([\[SNIA-CDMI\]](#)) from SNIA defines metadata that can be associated with data that is stored by a cloud storage provider. While the metadata currently defined do not match the needs of CDN Interconnection, it is worth considering CDMI as one of the existing pieces of work that may potentially be leveraged for the CDNI Metadata interface (e.g by extending the CDMI metadata to address more specific CDNI needs).

[C.2.](#) Related Research Projects

[C.2.1.](#) OCEAN

OCEAN (<http://www.ict-ocean.eu/>) is an EU funded research project that started in February 2010 for 3 years. Some of its objectives are relevant to CDNI. It aims, among other things, at designing a new architectural framework for audiovisual content delivery over the Internet, defining public interfaces between its major building blocks in order to foster multi-vendor solutions and interconnection between Content Networks (the term "Content Networks" corresponds here to the definition introduced in [\[RFC3466\]](#), which encompasses CDNs).

OCEAN has not yet published any open specifications, nor common best practices, defining how to achieve such CDN interconnection.

[C.2.2.](#) Eurescom P1955

Eurescom P1955 was a 2010 research project involving a four European Network operators, which studied the interests and feasibility of

interconnecting CDNs by firstly elaborating the main service models around CDN interconnection, as well as analyzing an adequate CDN interconnection technical architecture and framework, and finally by providing recommendations for telcos to implement CDN interconnection. The Eurescom P1955 project ended in July 2010.

The authors are not aware of material discussing CDN interconnection protocols or interfaces made publicly available as a deliverable of this project.

Authors' Addresses

Ben Niven-Jenkins
Velocix (Alcatel-Lucent)
326 Cambridge Science Park
Milton Road, Cambridge CB4 0WG
UK

Email: ben@velocix.com

Francois Le Faucheur
Cisco Systems
Greenside, 400 Avenue de Roumanille
Sophia Antipolis 06410
France

Phone: +33 4 97 23 26 19
Email: flefauch@cisco.com

Nabil Bitar
Verizon
40 Sylvan Road
Waltham, MA 02145
USA

Email: nabil.bitar@verizon.com

